

# Gene functional classification from heterogeneous data

Paul Pavlidis  
Columbia Genome Center  
Columbia University  
pp175@columbia.edu

Jason Weston  
Barnhill Technologies  
Savannah, Georgia  
jason@barnhilltechnologies.com

Jinsong Cai  
Department of Medical Informatics  
Columbia University  
jic7001@dmi.columbia.edu

William Noble Grundy\*  
Department of Computer Science  
Columbia University  
bgrundy@cs.columbia.edu

*Proceedings of the Fifth International Conference on Computational Molecular Biology,*  
April 21-24, 2001. pp. 242-248.

## Abstract

In our attempts to understand cellular function at the molecular level, we must be able to synthesize information from disparate types of genomic data. We consider the problem of inferring gene functional classifications from a heterogeneous data set consisting of DNA microarray expression measurements and phylogenetic profiles from whole-genome sequence comparisons. We demonstrate the application of the support vector machine (SVM) learning algorithm to this functional inference task. Our results suggest the importance of exploiting prior information about the heterogeneity of the data. In particular, we propose an SVM kernel function that is explicitly heterogeneous. We also show how to use knowledge about heterogeneity to aid in feature selection.

## 1 Introduction

A primary goal in biology is to understand the molecular machinery of the cell. The sequencing projects currently underway provide one view of this machinery. A complementary view is provided by data from DNA microarray hybridization experiments. In this paper, we describe computational techniques for inferring gene function from these two distinct types of data. These techniques are a first step toward the longer-term goal of learning about gene function simultaneously from many different types of genomic data.

Clearly, the availability of complete genomic sequence of human and other species provides a tremendous opportunity for understanding the functions of biological macromolecules. In this work, we infer gene function from phylogenetic profiles that are derived from a comparison between a given gene and a collection of complete genomes. Each profile characterizes the evolutionary history of a given gene. Two genes with similar phylogenetic profiles are likely to have similar functions, under the assumption that their similar pattern of inheritance across species is the result of a functional link.

Gene function can also be inferred from DNA microarray expression data. By offering a snapshot of the messenger RNA expression levels of thousands of genes at once, microarrays allow biologists to formulate models of gene expression on a scale that was unimaginable several years ago. Initial analyses of this type

---

\*Corresponding author: Department of Computer Science, Columbia University, 450 Computer Science Building, Mail Code 0401, 1214 Amsterdam Avenue, New York, NY 10027, Tel: (212) 939-7114, Fax: (212) 666-0140

of data focused on clustering algorithms, such as hierarchical clustering [10] and self-organizing maps [22]. These unsupervised algorithms attempt automatically to locate clusters of genes that share similar expression patterns and hence may share similarity in function. Recently, Brown *et al.* applied a collection of supervised learning techniques to a set of microarray expression data from yeast [3]. They showed that an algorithm known as a support vector machine (SVM) [23, 4, 7] provides excellent classification performance.

In this paper, we extend the methodology of Brown *et al.* to learn gene functional classifications from a heterogeneous data set consisting of microarray expression data and phylogenetic profiles. We show that this combination is most successful when the SVM operates in a feature space that is explicitly heterogeneous. SVMs are members of a larger class of algorithms, known as kernel methods, which can be non-linearly mapped to a higher-order feature space by replacing the dot product operation in the input space with a kernel function  $K(\cdot, \cdot)$ . Mercer's theorem [7] shows that every positive semi-definite kernel function corresponds to the dot product operation in some higher-dimensional feature space. In this work, we construct an explicitly heterogeneous kernel function by computing separate kernels for each data type and summing the results. The resulting kernel incorporates prior knowledge about the heterogeneity of the data by accounting for higher-order correlations among features of one data type but ignoring higher-order correlations across data types. This heterogeneous kernel leads to improved performance with respect to an SVM trained directly on the concatenated data.

We also show that prior knowledge of heterogeneity should be exploited when selecting subsets of input features for use in classification. For most of the gene functional classifications that we investigated, one type of genomic data provides significantly better training data than the other type. Many feature selection algorithms are available for automatically selecting the most useful features to use in training a classifier. We demonstrate that, for these data, feature selection algorithms that select among data types (i.e., learn from phylogenetic profiles, from gene expression data, or from both) perform better than algorithms that directly select features from the combined data set.

The idea of combining heterogeneous data sets to infer gene function is not new. Marcotte *et al.* describe an algorithm for functional annotation that uses expression vectors and phylogenetic profiles, as well as evolutionary evidence of domain fusion [14]. However, the algorithm consists of predicting functional links between pairs of genes using each type of data separately, and then cataloging the complete list of links. In contrast, the SVM method described here considers the various types of data at once, making a single prediction for each gene with respect to each functional category. Indeed, the performance of SVMs when data types are combined and a single hypothesis is formed is superior to combining two independent hypotheses, and we believe this will be true for a wide range of techniques.

## 2 Methods

The experiments carried out here use two types of genomic data. The first data set derives from a collection of DNA microarray hybridization experiments [10]. Each data point represents the logarithm of the ratio of expression levels of a particular gene under two different experimental conditions. The data consists of a set of 79-element gene expression vectors for 2465 yeast genes. These genes were selected by Eisen *et al.* based on the availability of accurate functional annotations. The data were generated from spotted arrays using samples collected at various time points during the diauxic shift [8], the mitotic cell division cycle [21], sporulation [6], and temperature and reducing shocks, and are available on the Stanford web site (<http://www-genome.stanford.edu>).

In addition to the microarray expression data, each of the 2465 yeast genes is characterized by a phylogenetic profile [17]. In its simplest form, a phylogenetic profile is a bit string, in which the Boolean value of each bit indicates whether the gene of interest has a close homolog in the corresponding genome. The profiles employed in this paper contain, at each position, the negative logarithm of the lowest  $E$ -value reported

by BLAST version 2.0 [1] in a search against a complete genome, with negative values (corresponding to E-values greater than 1) truncated to 0. Two genes in an organism can have similar phylogenetic profiles for one of two reasons. First, genes with a high level of sequence similarity will have, by definition, similar phylogenetic profiles. Second, for two genes which lack sequence similarity, the similarity in phylogenetic profiles reflects a similar pattern of occurrence of their homologs across species. This coupled inheritance may indicate a functional link between the genes, on the hypothesis that the genes are always present together or always both absent because they cannot function independently of one another. The profiles in this study are constructed using 24 complete genomes, collected from The Institute for Genomic Research website (<http://www.tigr.org/tdb>) or from the Sanger Centre website (<http://www.sanger.au.uk/>). Prior to learning, the gene expression and phylogenetic profile vectors are adjusted to have a mean of 0 and a variance of 1.

Classification experiments are carried out using gene functional categories from the Munich Information Center for Protein Sequences Yeast Genome Database (MYGD) (<http://www.mips.biochem.mpg.de/proj/yeast>). The database contains several hundred functional classes, whose definitions come from biochemical and genetic studies of gene function. The experiments reported here use classes containing ten or more genes, and which are not substantially encompassed by any other class used, amounting to 108 classes. The complete data set and corresponding classifications are available at <http://www.cs.columbia.edu/compbio>.

For each class, a support vector machine is trained to discriminate between class members and non-members. A support vector machine is a supervised learning algorithm developed over the past decade by Vapnik and others [23]. In the form employed here, SVMs learn binary classifications; i.e., the SVM learns to answer the question, “Does the given gene belong to functional class X?” where X is some category such as “ribosomal genes” or “sugar and carbohydrate transporters.” Support vector machines classify points by locating them with respect to a hyperplane that separates class members from non-class members in a high-dimensional feature space. The characteristics of the feature space are determined by a kernel function, which is selected *a priori*. The current experiments employ a kernel function that has been shown to produce good classification performance for some MYGD classes using this gene expression data set [3]. The function is a dot product raised to the third power:  $K(\vec{X}, \vec{Y}) = \left( (\vec{X} \cdot \vec{Y} / \sqrt{\vec{X} \cdot \vec{X}} \sqrt{\vec{Y} \cdot \vec{Y}}) + 1 \right)^3$ . This kernel function takes into account pairwise and tertiary correlations among gene expression measurements. The normalization term in the denominator projects the data onto the unit sphere. As in previous work, the SVM uses a soft margin that accounts for the disparity in the number of positive and negative examples for each class. For details about this adjustment, see [3]. A useful introduction to SVMs is available [7], as is the software used to perform these experiments (<http://www.cs.columbia.edu/compbio>).

The two types of data — gene expression and phylogenetic profiles — are combined in three different fashions, which we refer to as early, intermediate and late integration. These three methods are summarized in Figure 1. In early integration, the two types of vectors are concatenated to form a single set of length-103 vectors, which serve as input for the SVM. In intermediate integration, the kernel values for each type of data are pre-computed separately, and the resulting values are added together. These summed kernel values are used in the training of the SVM. Thus, given the above kernel function  $K(\cdot, \cdot)$ , the heterogeneous kernel is  $K(\vec{X}_g, \vec{Y}_g) + K(\vec{X}_p, \vec{Y}_p)$ , where subscripts denote gene expression and phylogenetic profile data, respectively. In late integration, one SVM is trained from each data type, and the resulting discriminant values are added together to produce a final discriminant for each gene.

In intermediate integration, the heterogeneous kernel we propose is an attempt to incorporate prior knowledge into the task at hand. The method creates local features that are polynomial relationships between inputs within a single type of data. These local features are then combined linearly to create global features. From these global features a hyperplane is constructed. In contrast to the feature space produced by the early integration method, polynomial relationships among different types of inputs are now ignored. This restriction reflects our intuition that the correlations of inputs within one type of data are stronger than

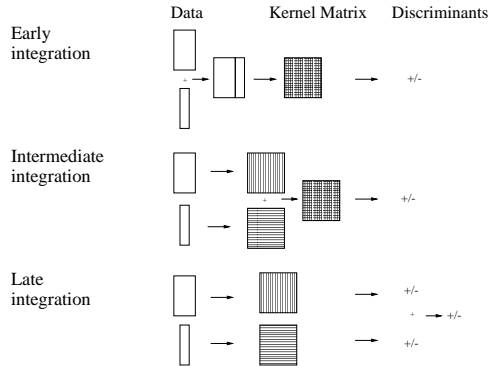


Figure 1: **Three methods for learning from heterogeneous data with a support vector machine.** In early integration, the two types of data are concatenated to form a single set of input vectors. In intermediate integration, the kernel values are computed separately for each data set and then summed. In late integration, one SVM is trained on each data type, and the resulting discriminant values are summed.

correlations between data types. In theoretical terms, removal of these correlations reduces overfitting as (unnecessary) capacity is reduced. Indeed, this approach is similar in spirit to one already used in digit recognition problems in order to incorporate prior knowledge about spatial location [23]. This incorporation was achieved by constructing sparse polynomials that sum across sub-kernels computed for many small patches within an image. In experiments on 400-pixel input spaces, the authors decreased the number of polynomial features from  $10^{23}$  to  $10^{14}$  and reported a 68% reduction in test error [23].

Each classification experiment is performed using cross-validation. For a given class, the positively labeled and negatively labeled genes are split randomly into  $n$  groups for  $n$ -fold cross-validation. An SVM is trained on  $n - 1$  of the groups and is tested on the remaining group. This procedure is repeated  $n$  times, each time using a different group of genes as a test set. For most of the experiments, we use three-fold cross-validation. Leave-one-out cross-validation, which is used in some of the experiments, is simply  $n$ -fold cross-validation with  $n$  equal to the total number of training examples.

The performance of each SVM is measured by examining how well the classifier identifies the positive and negative examples in the test sets. To judge overall performance, we define the cost of using the method  $M$  as  $C(M) = (fp(M) + 2 \cdot fn(M))/n$ , where  $fp(M)$  is the number of false positives for method  $M$ ,  $fn(M)$  is the number of false negatives for method  $M$ , and  $n$  is the number of members in the class. The false negatives are weighted more heavily than the false positives because, for these data, the number of positive examples is small compared to the number of negatives. To see why this unequal cost weighting is important, consider two classifiers A and B trained to recognize the class of histones, which contains 15 genes. Say that, on a test set of 822 genes, classifier A correctly identifies all 15 histones but also includes 15 non-histones in its list of positive genes. On the other hand, suppose that classifier B classifies everything in the test set as negative. Clearly, classifier A has learned to recognize something about the histones, whereas classifier B has learned nothing at all. However, using an equal weighting of false positive and false negatives, these two classifiers would yield the same cost (15). We assign a higher cost to false negatives in order to implement the intuition that failing to recognize one of the few positive examples is worse than inaccurately including one of the many negative examples in the test set. The cost for each method is compared to the cost  $C(N)$  for using the null learning procedure, which classifies all test examples as negative. We define the normalized cost savings of using the learning procedure  $M$  as  $S(M) = (C(N) - C(M))/2n$ , where  $n$  is the total number of positive examples in the class. Thus, a perfect classifier has a normalized cost savings of 1, and the null classifier has a normalized cost savings of 0.

We also performed feature selection on the combined data using the Fisher criterion score [2, 11]. For a given feature  $j$ , we compute the mean and standard deviation of that feature across the positive examples ( $\mu_j^+$  and  $\sigma_j^+$ , respectively) and across the negative examples ( $\mu_j^-$  and  $\sigma_j^-$ ). The Fisher criterion score,  $(\mu_j^+ - \mu_j^-)^2 / ((\sigma_j^+)^2 + (\sigma_j^-)^2)$ , gives higher values to features whose means differ greatly between the two classes, relative to their variances.

Finally, we also used the  $k$ -nearest neighbors algorithm (see, e.g [9]) as a means of selecting a type of data from which to learn. The algorithm labels a test point as positive if more than  $k/2$  of its closest (in Euclidean distance) neighbors from the training set are labeled positive; otherwise, the point is labeled negative. Taking  $k = 1$  (i.e., 1-nearest neighbor) amounts to assigning to each point the label associated with its closest neighbor. Given a set of features, one can measure the quality of the set via the leave-one-out error of this algorithm.

### 3 Results

Our initial experiments aimed at determining which classes among the 108 selected were learnable from either data type used alone. Based on the cost-savings measure, for each data type we selected the top 15 classes for further study. Three classes appear on both lists, yielding a total of 27 classes. The results are summarized in the first two columns of Table 1. Included in the table are the equivalents of all five classes used by Brown *et al.* [3].<sup>1</sup> These experiments show that the SVM methodology generalizes well to phylogenetic profile data, and that this new type of data allows for the characterization of new functional classes.

The second set of experiments tests the ability of the SVMs to learn from both types of data at once. The final three columns in Table 1 summarize these results, and Table 2 provides more details about the top five MYGD classes. Overall, integrating the data using a heterogeneous kernel function provides a normalized cost savings that is the best-performing or comparable to the best-performing method in 21 of the 27 classes, where “comparable” means that the two values differ by no more than the sum of their standard deviations. This is considerably more classes than any of the other four methods. Furthermore, the average cost savings across all classes is higher for this method than any other ( $p < 0.05$ , paired Student’s  $t$ -test). Similarly, the intermediate integration scheme fails to learn to classify only two classes, which is fewer classes than any of the other methods.

Learning from both data types is not always a good idea. For four classes, all three of the combined methods lead to decreased classification performance relative to an SVM trained on a single type of data. In every case, the decrease occurs when one data type provides much more information than the other, indicating that the inferior data type contributes noise that disrupts learning. This observation suggests that a feature selection algorithm that effectively eliminates noisy features should allow an SVM to learn these classes more accurately.

However, our experiments demonstrate that a naive feature selection algorithm, which does not take into account the heterogeneity of the data, does not typically yield improved classification performance. Figure 2 shows the results of using the Fisher criterion score to select features from the combined data. By treating each feature independently, this simple feature selection algorithm does not take into account possible correlations between features, but the algorithm has the advantages of simplicity and efficiency. For the most part, classification performance declines as features are removed. For several MYGD classes, when combining the two types of data leads to a substantial decline in performance, feature selection yields

---

<sup>1</sup>The MYGD functional catalog has been revised since the publication of Brown *et al.*, changing the composition of the classes substantially. For example, the “proteasome” class used in that paper has been subsumed into the “proteolysis” class. This is why the SVM performance on these classes differs from the performance reported in the earlier work.

Class	Exp	Phylo	Early	Intermediate	Late
amino acid transporters	0.05 ± 0.04	<b>0.77 ± 0.10</b>	0.50 ± 0.04	<b>0.71 ± 0.08</b>	0.49 ± 0.07
ribosomal proteins	0.71 ± 0.02	0.09 ± 0.03	<b>0.76 ± 0.01</b>	0.71 ± 0.01	0.69 ± 0.01
sugar and carbohydrate transporters	0.33 ± 0.07	<b>0.67 ± 0.02</b>	<b>0.68 ± 0.06</b>	<b>0.70 ± 0.01</b>	0.63 ± 0.03
glycolysis and gluconeogenesis	0.21 ± 0.03	<b>0.43 ± 0.05</b>	0.28 ± 0.02	<b>0.39 ± 0.05</b>	<b>0.39 ± 0.04</b>
mitochondrial organization	<b>0.40 ± 0.03</b>	0.15 ± 0.01	<b>0.43 ± 0.03</b>	<b>0.42 ± 0.02</b>	0.35 ± 0.02
tricarboxylic acid pathway	0.21 ± 0.11	0.15 ± 0.07	<b>0.32 ± 0.08</b>	<b>0.42 ± 0.07</b>	<b>0.25 ± 0.13</b>
deoxyribonucleotide metabolism	0.07 ± 0.05	<b>0.31 ± 0.11</b>	<b>0.24 ± 0.15</b>	<b>0.39 ± 0.11</b>	<b>0.31 ± 0.12</b>
organization of cytoplasm	0.35 ± 0.01	0.18 ± 0.01	<b>0.38 ± 0.01</b>	0.34 ± 0.02	<b>0.35 ± 0.02</b>
transport ATPases	0.13 ± 0.04	<b>0.37 ± 0.05</b>	0.23 ± 0.05	<b>0.32 ± 0.04</b>	0.22 ± 0.03
amino acid biosynthesis	0.18 ± 0.02	0.28 ± 0.02	0.29 ± 0.03	<b>0.36 ± 0.04</b>	0.27 ± 0.02
purine ribonucleotide metabolism	0.17 ± 0.03	<b>0.26 ± 0.05</b>	0.20 ± 0.04	<b>0.33 ± 0.04</b>	0.19 ± 0.03
pyrimidine ribonucleotide metabolism	0.03 ± 0.02	<b>0.33 ± 0.06</b>	0.11 ± 0.04	<b>0.28 ± 0.03</b>	0.17 ± 0.03
cytoplasmic degradation	<b>0.32 ± 0.01</b>		<b>0.32 ± 0.06</b>	<b>0.30 ± 0.03</b>	0.17 ± 0.02
respiration	<b>0.32 ± 0.02</b>		<b>0.30 ± 0.04</b>	0.23 ± 0.04	0.17 ± 0.03
organization of chromosome structure	<b>0.31 ± 0.01</b>		<b>0.30 ± 0.01</b>	<b>0.29 ± 0.02</b>	0.13 ± 0.03
phosphate utilization	<b>0.22 ± 0.04</b>	0.08 ± 0.05	<b>0.26 ± 0.05</b>	<b>0.21 ± 0.04</b>	<b>0.22 ± 0.04</b>
organization of plasma membrane	0.07 ± 0.02	<b>0.25 ± 0.01</b>	<b>0.24 ± 0.03</b>	<b>0.26 ± 0.03</b>	<b>0.26 ± 0.02</b>
pentose phosphate pathway		<b>0.20 ± 0.15</b>		<b>0.26 ± 0.07</b>	<b>0.15 ± 0.10</b>
cellular import	0.04 ± 0.02	<b>0.25 ± 0.04</b>	<b>0.18 ± 0.05</b>	0.17 ± 0.03	<b>0.21 ± 0.04</b>
protein folding and stabilization		<b>0.24 ± 0.04</b>	<b>0.20 ± 0.04</b>	<b>0.23 ± 0.05</b>	0.14 ± 0.04
proteolysis	<b>0.23 ± 0.02</b>		<b>0.24 ± 0.02</b>	<b>0.18 ± 0.06</b>	0.17 ± 0.01
pheromone response generation	<b>0.24 ± 0.05</b>		0.15 ± 0.03	<b>0.14 ± 0.08</b>	
nuclear organization	<b>0.21 ± 0.01</b>	0.07 ± 0.01	<b>0.24 ± 0.03</b>	<b>0.24 ± 0.02</b>	0.17 ± 0.02
drug transporters		<b>0.23 ± 0.09</b>			
organization of endoplasmatic reticulum	<b>0.20 ± 0.02</b>		<b>0.22 ± 0.03</b>	<b>0.19 ± 0.05</b>	0.13 ± 0.03
organization of cell wall	<b>0.12 ± 0.04</b>	<b>0.19 ± 0.06</b>	<b>0.14 ± 0.08</b>	<b>0.16 ± 0.07</b>	<b>0.21 ± 0.08</b>
anion transporters		<b>0.21 ± 0.02</b>			
Mean cost savings	0.19 ± 0.02	0.21 ± 0.04	0.27 ± 0.03	0.31 ± 0.03	0.24 ± 0.03
Number of best-performing	10	12	17	21	8
Number of non-learnable	4	6	3	2	3

Table 1: **Summary of learning performance results.** Each row in the table contains the cost savings for one MYGD classification. Each cost savings is computed via three-fold cross-validation, with standard deviation calculated across five repetitions. The first two columns are from SVMs trained on a single type of data (gene expression or phylogenetic profiles). The remaining three columns are from SVMs trained using early, intermediate and late integration of the data, as described in the text. Values in bold face are the best-performing or are comparable to the best-performing of the five methods. A missing value indicates that the cost savings is not significantly greater than zero. The last three rows are summary statistics, giving the average values for each method and total number of bold-face and missing values in each column.

Class	Size	FP	FN
amino acid transporters	22	2.0 ± 0.4	5.6 ± 0.2
ribosomal proteins	173	26.6 ± 1.2	34.2 ± 1.1
sugar and carbohydrate transporters	32	2.4 ± 0.7	9.0 ± 0.0
deoxyribonucleotide metabolism	9	0.2 ± 0.2	4.6 ± 0.7
mitochondrial organization	296	84.8 ± 1.8	128.4 ± 1.7

Table 2: **Error rates for selected MYGD classes.** The table lists error rates for the five most learnable MYGD classes. Each row contains the name and size of the class, and the average numbers of false positives and false negatives for that class from an SVM using intermediate integration.




Figure 2: **Effect of feature selection on learning performance.** Each series shows the performance of an SVM trained on the combined data set using varying numbers of features, selected according to the Fisher criterion score. These examples are representative of results obtained for all 27 classes tested.

	Express	Phylo	Early	Fisher sel	1-NN sel	5-CV sel
wins in all 3 trials	8	15	17	19	20	20
wins in 2 or more trials	8	15	17	22	25	26

Table 3: **Feature selection for optimal data set choice.** The table lists the number of classes for which various data set selection algorithms choose the best or within one standard deviation of the best performing data set. The results are computed over three separate trials, so the values are given for the case when the best choice is made in all three trials or in more than two trials. The algorithms (columns from left to right) are: always choosing the expression, phylogenetic or combined data, and selecting the dataset via the Fisher criterion score, leave-one-out error of the 1-nearest neighbor algorithm and five-fold cross-validation of SVMs.

a small improvement. However, in none of these cases does feature selection yield a level of performance comparable to that obtained using a single data type.

A more sophisticated feature selection method, which does take into account correlations and nonlinearities, yields similar results. The method uses the SVM solution to measure the quality of the features, and removes features that appear to contribute least — a so-called filter method (see, e.g [24]). We do not give details of the method here because, although the results are marginally superior to the Fisher criterion results, the same problem arises: the SVM feature selection method does not achieve performance equal to the best-performing single data type in cases where the combined data set performs poorly. Even worse, for these cases, hand-picking the best-performing single data set and trying various ways of adding features to it still leads to a consistent deterioration in performance. Apparently, when one data set performs quite poorly compared to the other (e.g., amino acid transporters, glycolysis and gluconeogenesis, sugar and carbohydrate transporters, etc.), no useful information can be gleaned from it.

The problem that the combined data sometimes performs poorly (or in general one dataset/kernel combination performs better than others) suggests that it would be useful to determine that such an outcome was likely before attempting to train. In other words, one would like to be able to make the optimal choice between the options of using one data type, or the combination, before training.

Because this task is a special case of feature selection (selecting the best set of features given  $n$  distinct sets), traditional quality measures for feature selection can still be used. For example, the Fisher criterion

score can be used to select the best data set by choosing the largest of the mean Fisher criterion scores.<sup>2</sup> One can also attempt to estimate the SVM’s performance directly, via generalization bounds such as the VC bounds [23] or the span bound [5], or via cross-validation. We found that only the cross-validation was useful: the VC bound and span bound are too loose in this case. Finally, we also measure the leave-one-out error of the nearest neighbor algorithm, which can be considered as a compromise between the Fisher criterion score and cross-validation in terms of computation speed versus accuracy.

We performed a first test of this idea by selecting among three data sets: the gene expression data, phylogenetic profiles or concatenated data set. For the 27 classes, we counted how many times in three trials each method chose the best-performing data set. These results, shown in Table 3, indicate that choosing the correct data set ahead of time can give improved results. Both cross-validation and nearest neighbors are superior to the Fisher criterion score, since even when these methods do not choose the outright best data set they choose one close to the best (as hinted by the “wins in 2 or more trials” results). However, while cross-validation gains this improvement at the expense of high computational cost, nearest neighbors provides a cheap compromise. Bearing in mind that nearest neighbors can work in a feature space via kernels [7], these preliminary results suggest that one can estimate the performance of any kernel (including the intermediate integration method) with fairly low computational cost.

## 4 Discussion

As the quantity and variety of genomic data increases, molecular biology shifts from a hypothesis-driven model to a data-driven one. Whereas previously a single laboratory could collect data and test hypotheses regarding a single system or pathway, this new paradigm requires combining genome-wide experimental results, typically gathered and shared across multiple laboratories. For example, constructing a single,  $n$ -element phylogenetic profile requires the availability of  $n$  complete genomic sequences, which clearly could not yet be generated by a single laboratory. The data-driven model requires sophisticated computational techniques that handle very large, heterogeneous data sets.

The support vector machine learning algorithm is such a technique. SVMs scale well and have been used successfully with large training sets in the domains of text categorization and image recognition [7]. Furthermore, in this paper, we demonstrate that SVMs can learn from heterogeneous data sets. With an appropriate kernel function, the SVM learns from a combination of two different types of feature vectors. In most cases, the resulting trained SVM provides as good or better gene functional classification performance than an SVM trained on either data set alone.

For these data and these classifications, the heterogeneous kernel we introduce here (the intermediate integration method) performs better than the other techniques we investigated. We hypothesize that the improved performance results from the kernel’s ability to exploit our prior knowledge that correlations within one type of data are stronger than correlations between different types. However, it may be useful to consider feature selection algorithms that determine which type of data is most appropriate for a given class, because no single method provides the consistently best performance across all classes. Data type selection can be addressed by using a simple feature selection metric, such as the Fisher criterion score, or more reliably (but with more computational expense) via cross-validation. Using the latter, from preliminary experiments we expect one to obtain near best performance across all classes.

Our results show that the supervised learning methodology proposed by Brown *et al.* [3] can be extended in a straightforward fashion to some other classes in the MYGD. We have also shown, however, that the majority of the MYGD classifications are not learnable from either gene expression data or phylogenetic profiles. We do not believe that the failure to learn many of these classes is a failure of the SVM method.

---

<sup>2</sup>The mean Fisher score of the combined data set is always intermediate between the mean Fisher scores from the two subsets. Hence, we choose the combined data set if its score falls within 10% of the highest alternative.



Rather, for many functional classes, the data are simply not informative. The expression data is only informative if the genes in the class are coordinately regulated at the level of transcription under the conditions tested. Similarly, phylogenetic profiles are limited in resolution in part because relatively few genomes are available. In particular, among the genomes from which we derived the phylogenetic profiles, all but one are bacterial. Thus it is difficult to generate useful phylogenetic profiles for genes that are specific to eucaryotes. Because the availability of expression, sequence, and other kinds of data are increasing steadily, we expect that the tools we are developing will continue to improve in power.

In our experiments, the primary utility of the phylogenetic profiles appears to lie in their ability to summarize sequence similarity information rather than the inheritance patterns of genes during various speciation events. Analysis of the classes that are most easily learnable from phylogenetic profiles alone shows that these are also the classes that share considerable sequence similarity among their members, for example, the various transporter classes (data not shown). In a previously published report on the use of phylogenetic profiles in yeast [14], this effect was eliminated by merging groups of similar genes, and by making links between pairs of genes, rather than requiring that the phylogenetic profile similarity extend throughout an entire functional class, as we have. In our experiments, removing or combining the data for genes with sequence similarity would have had the undesirable effect of forcing the combining of the corresponding expression data, and there is no reason to think that genes with sequence similarity would generally be coordinately regulated at the expression level. However, there are obviously benefits to considering sequence similarity in a gene classification task, and in the future we will consider other techniques for summarizing sequence similarities in a fixed-length vector.

In the experiments reported here, we have used a fixed kernel function for all of the SVM learning in order to allow direct comparisons across different feature combination and feature selection algorithms. We selected this particular kernel because of its straightforward interpretation (accounting for all primary features and up to tertiary correlations in the data) and because previous work showed that this kernel performs well for this gene expression data set [3]. However, we did not experiment with the radial basis kernel, which also gave good performance previously, and we did not optimize the kernel with respect to the phylogenetic profile data. The performance of the various algorithms would no doubt be improved by empirical kernel optimization. There is no reason to suppose, however, that a particular method would benefit more than others from such optimization.

The experiments reported here suggest several avenues for future research. One obvious research direction involves including additional types of data. Having shown that two types of data can be fruitfully combined, we plan to extend the techniques described here to feature vectors derived from, for example, the upstream promoter regions of genes [16]. We also plan to experiment with the Fisher kernel method, in which each type of data is compared to a probabilistic model of the domain [13, 12]. By converting the heterogeneous features to probability gradients, we hope to make the various types of data more directly comparable.

Support vector machines are part of a larger class of algorithms known as kernel methods, which have recently been gaining in popularity [18]. A kernel method is any algorithm that employs a kernel function to implicitly operate in a higher-dimensional space. In addition to SVM classifiers, kernel methods have been developed for regression [19], discriminant analysis [15], and principal components analysis [20]. More members of this promising class of algorithms should be applied to problems in computational biology.

### **Acknowledgments**

The authors wish to thank Ilya Muchnik and Vladimir Vapnik for helpful discussions. WNG is funded by an Award in Bioinformatics from the PhRMA Foundation, and by National Science Foundation grant DBI-0078523.

## References

- [1] S. F. Altschul, T. L. Madden, A. A. Schaffer, J. Zhang, Z. Zhang, W. Miller, and D. J. Lipman. Gapped BLAST and PSI-BLAST: A new generation of protein database search programs. *Nucleic Acids Research*, 25:3389–3402, 1997.
- [2] C. Bishop. *Neural Networks for Pattern Recognition*. Oxford UP, Oxford, UK, 1995.
- [3] M. P. S. Brown, W. N. Grundy, D. Lin, N. Cristianini, C. Sugnet, T. S. Furey, Jr. M. Ares, and D. Haussler. Knowledge-based analysis of microarray gene expression data using support vector machines. *Proceedings of the National Academy of Sciences of the United States of America*, 97(1):262–267, 2000.
- [4] C. J. C. Burges. A tutorial on support vector machines for pattern recognition. *Data Mining and Knowledge Discovery*, 2(2):121–167, 1998.
- [5] O. Chapelle and V. Vapnik. Model selection for support vector machines. In Sara A. Solla, Todd K. Leen, and Klaus-Robert Müller, editors, *Advances in Neural Information Processing Systems 12*. MIT Press, 2000.
- [6] S. Chu, J. DeRisi, M. Eisen, J. Mulholland, D. Botstein, P. Brown, and I. Herskowitz. The transcriptional program of sporulation in budding yeast. *Science*, 282:699–705, 1998.
- [7] N. Cristianini and J. Shawe-Taylor. *An Introduction to Support Vector Machines*. Cambridge UP, 2000.
- [8] J.L. DeRisi, V.R. Iyer, and P.O. Brown. Exploring the metabolic and genetic control of gene expression on a genomic scale. *Science*, 278:680–686, 1997.
- [9] R. O. Duda and P. E. Hart. *Pattern Classification and Scene Analysis*. Wiley, New York, 1973.
- [10] M. Eisen, P. Spellman, P.O. Brown, and D. Botstein. Cluster analysis and display of genome-wide expression patterns. *Proceedings of the National Academy of Sciences of the United States of America*, 95:14863–14868, 1998.
- [11] T. S. Furey, N. Cristianini, N. Duffy, D. W. Bednarski, M. Schummer, and D. Haussler. Support vector machine classification and validation of cancer tissue samples using microarray expression data. *Bioinformatics*, 2001. To appear.
- [12] T. Jaakkola, M. Diekhans, and D. Haussler. Using the Fisher kernel method to detect remote protein homologies. In *Proceedings of the Seventh International Conference on Intelligent Systems for Molecular Biology*, pages 149–158, Menlo Park, CA, 1999. AAAI Press.
- [13] T. Jaakkola and D. Haussler. Exploiting generative models in discriminative classifiers. In *Advances in Neural Information Processing Systems 11*, San Mateo, CA, 1998. Morgan Kaufmann.
- [14] E. M. Marcotte, M. Pellegrini, M. J. Thompson, T. O. Yeates, and D. Eisenberg. A combined algorithm for genome-wide prediction of protein function. *Nature*, 402(6757):83–86, 1999.
- [15] S. Mika, G. Rätsch, J. Weston, B. Schölkopf, and K.-R. Müller. Fisher discriminant analysis with kernels. In *Proceedings of the IEEE Neural Networks for Signal Processing Workshop 1999*, 1999.
- [16] P. Pavlidis, T. S. Furey, M. Liberto, D. Haussler, and W. N. Grundy. Promoter region-based classification of genes. In *Proceedings of the Pacific Symposium on Biocomputing*, 2001. To appear.
- [17] M. Pellegrini, E. M. Marcotte, M. J. Thompson, D. Eisenberg, and T. O. Yeates. Assigning protein functions by comparative genome analysis: protein phylogenetic profiles. *Proceedings of the National Academy of Sciences of the United States of America*, 96(8):4285–4288, 1999.
- [18] B. Schölkopf, C. J. C. Burges, and A. J. Smola, editors. *Advances in Kernel Methods: Support Vector Learning*. MIT Press, Cambridge, MA, 1999.
- [19] B. Schölkopf, A. Smola, and K.-R. Müller. Nonlinear component analysis as a kernel eigenvalue problem. *Neural Computation*, 10(5):1299–1319, 1996.
- [20] B. Schölkopf, A. Smola, and K.-R. Müller. Kernel principal component analysis. In *Proceedings ICANN97*, Springer Lecture Notes in Computer Science, page 583, 1997.

- [21] P. T. Spellman, G. Sherlock, M. Q. Zhang, V. R. Iyer, K. Anders, M. B. Eisen, P. O. Brown, D. Botstein, and B. Futcher. Comprehensive identification of cell cycle-regulated genes of the yeast *Saccharomyces cerevisiae* by microarray hybridization. *Mol Biol Cell*, 9:3273–3297, 1998.
- [22] P. Tamayo, D. Slonim, J. Mesirov, Q. Zhu, S. Kitareewan, E. Dmitrovsky, E. Lander, and T. Golub. Interpreting patterns of gene expression with self-organizing maps. *Proceedings of the National Academy of Sciences of the United States of America*, 96:2907–2912, 1999.
- [23] V. N. Vapnik. *Statistical Learning Theory*. Adaptive and learning systems for signal processing, communications, and control. Wiley, New York, 1998.
- [24] J. Weston, S. Mukherjee, O. Chapelle, M. Pontil, T. Poggio, and V. Vapnik. Feature selection for SVMs. In Sara A Solla, Todd K Leen, and Klaus-Robert Müller, editors, *Advances in Neural Information Processing Systems 13*. MIT Press, 2001.