

# A Comprehensive Evaluation of Multicategory Classification Methods for Microarray Gene Expression Cancer Diagnosis

Alexander Statnikov<sup>1</sup>, Constantin F. Aliferis<sup>1</sup>, Ioannis Tsamardinos<sup>1</sup>, Douglas Hardin<sup>2</sup>, Shawn Levy<sup>1</sup>

<sup>1</sup>Department of Biomedical Informatics, <sup>2</sup>Department of Mathematics, Vanderbilt University, Nashville, TN, USA

## Abstract

**Motivation:** Cancer diagnosis is one of the most important emerging clinical applications of gene expression microarray technology. We are seeking to develop a computer system for powerful and reliable cancer diagnostic model creation based on microarray data. To keep a realistic perspective on clinical applications we focus on multicategory diagnosis. In order to equip the system with the optimum combination of classifier, gene selection and cross-validation methods, we performed a systematic and comprehensive evaluation of several major algorithms for multicategory classification, several gene selection methods, multiple ensemble classifier methods, and two cross validation designs using 11 datasets spanning 74 diagnostic categories and 41 cancer types and 12 normal tissue types.

**Results:** Multicategory Support Vector Machines (MC-SVMs) are the most effective classifiers in performing accurate cancer diagnosis from gene expression data. The MC-SVM techniques by Crammer and Singer, Weston and Watkins, and one-versus-rest were found to be the best methods in this domain. MC-SVMs outperform other popular machine learning algorithms such as K-Nearest Neighbors, Backpropagation and Probabilistic Neural Networks, often to a remarkable degree. Gene selection techniques can significantly improve classification performance of both MC-SVMs and other non-SVM learning algorithms. Ensemble classifiers do not generally improve performance of the best non-ensemble models. These results guided the construction of a *software system GEMS (Gene Expression Model Selector) that automates high-quality model construction and enforces sound optimization and performance estimation procedures*. This is the first such system to be informed by a rigorous comparative analysis of the available algorithms and datasets.

**Availability:** The software system *GEMS* is available for download from <http://www.gems-system.org> for non-commercial use.

**Contact:** [alexander.statnikov@vanderbilt.edu](mailto:alexander.statnikov@vanderbilt.edu)

## 1. Introduction

An important emerging medical application domain for microarray gene expression profiling technology is clinical decision support in the form of diagnosis of disease as well as prediction of clinical outcomes in response to treatment. The two areas in medicine that currently attract the greatest attention in that respect are management of cancer and infectious diseases ([Fortina2002] and [Ntzani2003]).

A necessary prerequisite for the creation of clinically successful microarray-based diagnostic models is a solid understanding of the relative strengths and weaknesses of available classification and related methods (i.e., gene selection and cross-validation). While prior research has established the feasibility of creating accurate models for cancer diagnosis, the corresponding studies conducted limited experiments in terms of the number of classifiers, gene selection algorithms, number of datasets, and types of cancer involved (e.g., [Yeo2001], [Su2001], [Ramaswamy2001], [Yeang2001], and [Lee2003]). In addition, the results of these studies cannot be combined into a comprehensive comparative meta-analysis because each study follows different experimental protocols and applies learning algorithms differently. Thus, it is not clear from the literature which classifier, if any, performs best among the many available alternatives. It is also currently poorly understood what are the best combinations of classification and gene selection algorithms across most array-based cancer datasets.

Another major methodological concern is the problem of *overfitting*, that is creating diagnostic models that may not generalize well to new data from the same cancer types and data distribution despite excellent performance on the training set. Since many algorithms are highly parametric and datasets consist of a relatively small number of high-dimensional samples, it is easy to overfit both the classifiers and the gene selection procedures especially when using intensive model search and powerful learners. Indeed recently, a number of reports appeared in the literature raising doubts about the generalization ability of classifiers produced by major studies in the field ([Schwarzer2000], [Reunanen2003], and [Guyon2003]). In recent meta-analytic assessment of 84 published microarray cancer outcome predictive studies [Ntzani2003], it was found that 74% of studies did not perform independent validation or cross-validation of proposed findings, 13% applied cross-validation in an incomplete fashion, and only 13% performed cross-validation correctly. On the other hand, when building a diagnostic model, one should also avoid *underfitting*, which results in classifiers that are not optimally robust due to limited experimentation. In particular, this is manifested by application of a specific learning algorithm without consideration of alternatives, or use of parametric learners with certain values of parameters without searching for the best ones.

At the heart of the present work lies the thesis that *an automated system can help with the creation of high-quality and robust diagnostic and prognostic models*. For such a system to be successful it must implement the best possible classification and gene selection algorithms for the domain and guide model selection by enforcing sound principles of model building and data analysis. Hence, to inform the development of such a system, the goals of the present work are to: (a) investigate which

one among the many powerful classifiers currently available for gene expression diagnosis performs the best *across many cancer types*; (b) how classifiers interact with existing gene selection methods in datasets with varying sample size, number of genes, and cancer types; (c) whether it is possible to increase diagnostic performance further using meta-learning in the form of ensemble classification; and (d) how to parameterize the classifiers and gene selection procedures so as to *avoid overfitting*. The ultimate goal is to utilize the knowledge gleaned from the above experiments to create a fully-automated software platform that creates high-quality, if not optimal, diagnostic models for use in clinical applications. A first incarnation of such a system is introduced as a result of the reported experiments.

## 2. Methods and Materials

### 2.1. Support Vector Machine-based classification methods

Support Vector Machines (SVMs) [Vapnik1998] are arguably the single most important development in supervised classification of recent years. SVMs often achieve superior classification performance compared to other learning algorithms across most domains and tasks; they are fairly insensitive to the *curse of dimensionality* and are efficient enough to handle very large-scale classification in both sample and variables. In clinical bioinformatics they have allowed construction of powerful experimental cancer diagnostic models based on gene expression data with thousands of variables and as little as few dozens samples (e.g., [Furey2000], [Guyon2002], and [Aliferis2003a]). Moreover, several efficient and high-quality implementations of SVM algorithms (e.g., [Joachims1999] and [Chang2003]) facilitate application of these techniques in practice. The first generation of SVMs could only be applied to binary classification tasks. Yet, most real-life diagnostic tasks are not binary. Moreover, all other things being equal, multicategory classification is significantly harder than binary classification [Mukherjee2003]. Fortunately, several algorithms have emerged during the last few years that allow multicategory classification with SVMs. The preliminary experimental evidence currently available suggests that some multicategory SVMs (MC-SVMs) perform well in isolated gene expression-based cancer diagnostic experiments ([Yeo2001], [Su2001], [Ramaswamy2001], [Yeang2001], and [Lee2003]).

Below we outline the principles behind SVM algorithms used in the study. Full technical descriptions can be found in the references provided in text. A detailed review of binary SVMs, exact mathematical formulations of both binary and multiclass SVM algorithms, and an illustration of MC-SVMs methods via a solution of example cancer diagnostic problem are presented in Appendices A, B, and C, respectively, which are available online from [Statnikov2004]. In the description of methods below,  $k$  is the number of classes or distinct diagnostic categories, and  $n$  is the number of samples or patients in the training dataset.

**Binary SVMs.** The main idea of binary SVMs is to implicitly map data to a higher dimensional space via a kernel function and then solve an optimization problem to identify the maximum-margin hyperplane that separates training instances [Vapnik1998]. The hyperplane is based on a set of boundary training instances, called *support vectors*. New instances are classified according to the side of the hyperplane they fall into (**Figure 1**). The optimization problem is most often formulated in a way that allows for non-separable data by penalizing misclassifications.

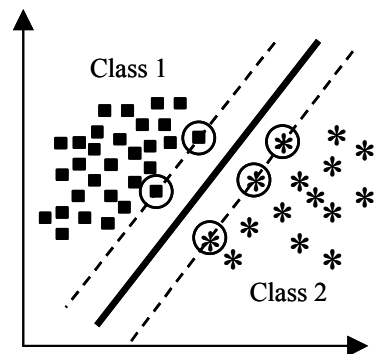
**Multiclass SVMs: One-Versus-Rest (OVR).** This is conceptually the simplest multiclass SVM method (see [Kressel1999] for details). Here we construct  $k$  binary SVM classifiers: class 1 (positive) versus all other classes (negative), class 2 versus all other classes, ..., class  $k$  versus all other classes (**Figure 2a**). The combined OVR decision function chooses the class of a sample that corresponds to the maximum value of  $k$  binary decision functions specified by the furthest “positive” hyperplane. By doing so, the decision hyperplanes calculated by  $k$  SVMs “shift”, which questions the optimality of the multicategory classification.

This approach is computationally expensive, since we need to solve  $k$  quadratic programming (QP) optimization problems of size  $n$ . Moreover, this technique does not currently have theoretical justification such as the analysis of generalization, which is a relevant property of a robust learning algorithm.

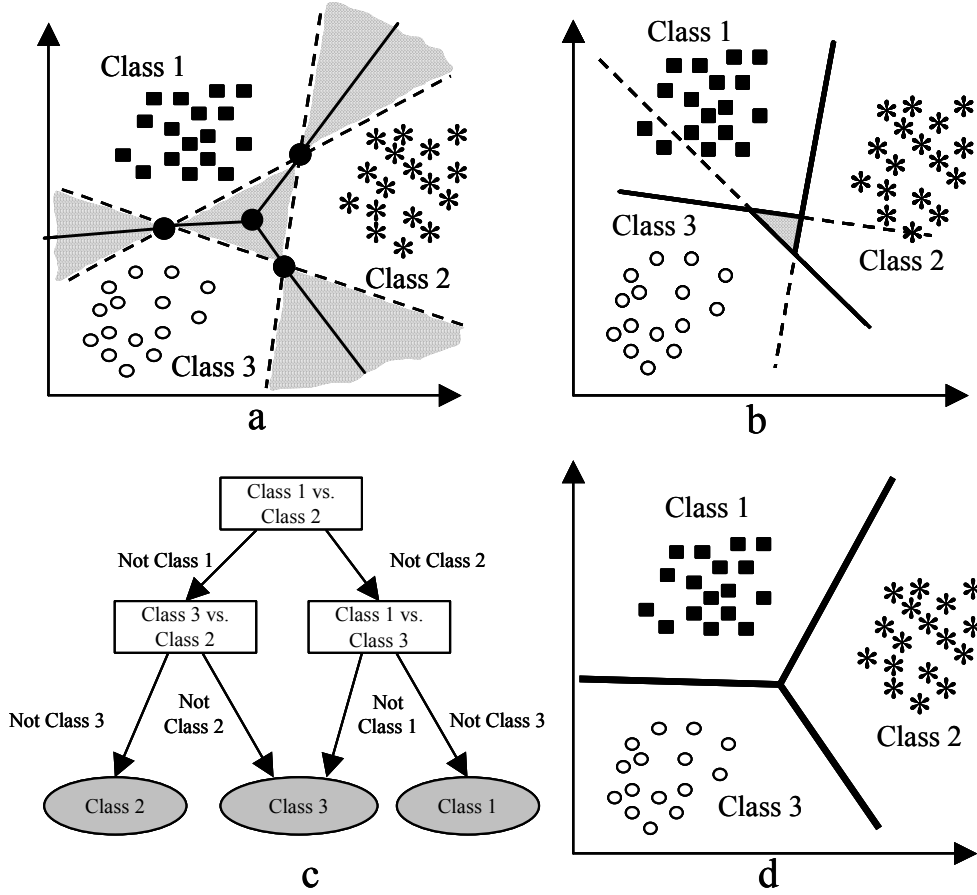
**Multiclass SVMs: One-Versus-One (OVO).** This method involves construction of binary SVM classifiers for all pairs of classes; in total there are  $\binom{k}{2} = \frac{k(k-1)}{2}$  pairs

(see **Figure 2b** and [Kressel1999]). In other words, for every pair of classes, a binary SVM problem is solved (with the underlying optimization problem to maximize the margin between two classes). The decision function assigns an instance to a class which has the largest number of votes, so-called *Max Wins strategy* [Friedman1996]. If ties still occur, a sample will be assigned a label based on the classification provided by the furthest hyperplane.

One of the benefits of this approach is that for every pair of classes we deal with a much smaller optimization problem, and in total we need to solve  $k(k-1)/2$  QP problems of size *smaller than*  $n$ . Given that QP optimization algorithms used for SVMs



**Figure 1.** A binary SVM selects a hyperplane (bold line) that maximizes the width of the “gap” (margin) between the two classes. The hyperplane is specified by “boundary” training instances, called support vectors shown with circles. New cases are classified according to the side of the hyperplane they fall into.



**Figure 2.** MC-SVM algorithms applied to a three-class diagnostic problem. (a) MC-SVM One-Versus-Rest constructs 3 classifiers: (1) class 1 vs classes 2 and 3, (2) 2 vs 1 and 3, and (3) 3 vs 1 and 2. (b) MC-SVM One-Versus-One constructs 3 classifiers: (1) class 1 vs class 2, (2) 2 vs 3, and (3) 1 vs 3. (c) MC-SVM DAGSVM constructs a decision tree on the basis of One-Versus-One SVM classifiers. (d) MC-SVM methods by Weston and Watkins and by Crammer and Singer construct a single classifier by maximizing margin between all classes simultaneously.

are polynomial to the problem size, such a reduction can yield substantial savings in the total computational time. Moreover, some researchers postulate that even if the entire multicategory problem is non-separable, while some of the binary sub-problems are separable, then OVO can lead to improvement of classification compared to OVR [Kressel1999]. Unlike the OVR approach, here tie-breaking plays only a minor role and does not affect the decision boundaries significantly. On the other hand, similarly to OVR, OVO does not currently have established bounds on the generalization error.

**Multiclass SVMs: DAGSVM.** The training phase of this algorithm is similar to the OVO approach using multiple binary SVM classifiers; however the testing phase of DAGSVM requires construction of a rooted binary decision directed acyclic graph (DDAG) using  $\binom{k}{2}$  classifiers (see **Figure 2c** and [Platt2000]). Each node of this graph is a binary SVM for a pair of

classes, say  $(p, q)$ . On the topologically lowest level there are  $k$  leaves corresponding to  $k$  classification decisions. Every non-leaf node  $(p, q)$  has two edges – the left edge corresponds to decision “not  $p$ ” and the right one corresponds to “not  $q$ ”. The choice of the class order in the DDAG list can be arbitrary as shown empirically in [Platt2000].

In addition to inherited advantages from the OVO method, DAGSVM is characterized by a bound on the generalization error.

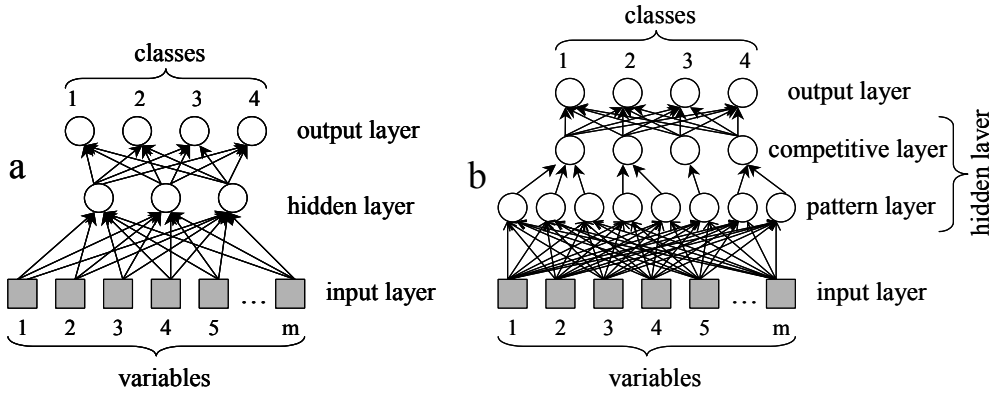
**Multiclass SVMs: Method by Weston and Watkins (WW).** This approach to multiclass SVMs is viewed by some researchers as a natural extension of the binary SVM classification problem (see **Figure 2d**, [Hsu2002] and [Weston1999]). Here, in the  $k$ -class case one has to solve a single quadratic optimization problem of size  $(k-1)n$  which is identical to binary SVMs for the case  $k=2$ . In a slightly different formulation of QP problem, a bounded formulation, decomposition techniques can provide a significant speed-up in the solution of the optimization problem ([Hsu2002], [Platt1999]). This method does not have an established bound on the generalization error, and its optimality is not currently proved.

**Multiclass SVMs: Method by Crammer and Singer (CS).** This technique is similar to WW (see **Figure 2d**, [Hsu2002] and [Crammer2000]). It requires solution of a single QP problem of size  $(k-1)n$ , however uses less slack variables in the constraints of the optimization problem, and hence it is cheaper computationally. Similarly to WW, the use of decompositions can provide a significant speed-up in the solution of the optimization problem [Hsu2002]. Unfortunately, the optimality of CS, as well as the bounds on generalization has not been demonstrated yet.

## 2.2. Non-SVM classification methods

In addition to five MC-SVM methods, three popular classifiers, K-Nearest Neighbors (KNN), Backpropagation Neural Networks (NN), and Probabilistic Neural Networks (PNN) were also used in this study. These learning methods have been exten-

**Figure 3.** Simplified illustration of the design of neural networks for a 4-category diagnostic problem with  $m$ -dimensional samples of variables (genes) and training set containing  $N$  samples. **(a)** Backpropagation Neural Network contains inputs for  $m$  variables (genes); hidden layer contains 3 units (this number is usually determined by cross-validation); and output layer contains a unit for each diagnostic category (1-of- $n$  encoding scheme). **(b)** Probabilistic Neural Network contains inputs for  $m$  variables (genes); pattern layer contains  $N$  units (a unit for each training instance); competitive layer contains 4 units (a unit for each diagnostic category) and receive inputs only from pattern units that are associated with the category to which the training instance belongs; and output layer contains a unit for each diagnostic category.



sively and successfully applied to gene expression based cancer diagnosis (e.g., [Khan2001], [Ramaswamy2001], [Pomeroy2002], [Nutt2003], [Singh2002], and [Berrar2003]).

**K-Nearest Neighbors (KNN).** The main idea of KNN is that it treats all samples as points in the  $m$ -dimensional space (where  $m$  is the number of variables) and given an unseen sample  $x$ , the algorithm classifies it by a vote of  $K$  nearest training instances as determined by some distance metric, typically Euclidian distance [Mitchell1997].

**Backpropagation Neural Networks (NN).** Backpropagation Neural Networks are feed-forward neural networks with signals propagated only forward through the layers of units. These networks are composed of (I) an input layer of units, which we feed with gene expression data, (II) hidden layer(s) of units, and (III) an output layer of units, one for each diagnostic category, so-called *1-of- $n$  encoding* (see **Figure 3a** and [Mitchell1997]). The connections among units have weights and are adjusted during the training phase (epochs of a neural network) by backpropagation learning algorithm. This algorithm adjusts weights by propagating the error between network outputs and true diagnoses backward through the network and employs gradient descent optimization to minimize the error function. This process is repeated until we find a vector of weights that best fits the training data. When training of a neural network is complete, unseen data instances are fed to the input units, propagated forward through the network, and the network outputs classifications.

**Probabilistic Neural Networks (PNN).** Probabilistic Neural Networks belong to the family of Radial Basis Function (RBF) neural networks [Mitchell1997]. RBF networks are feed-forward neural networks with only one hidden layer. The primary difference between a backpropagation neural network with one hidden layer and an RBF network is that for the latter one, the inputs are passed directly to the hidden layer *without weights*. The Gaussian density function is used in a hidden layer as an activation function. The weights for the connections among the hidden and the output layer are optimized via a least squares optimization algorithm. A key advantage of RBF networks is that they are trained much more efficiently than backpropagation neural networks.

Probabilistic Neural Networks are made up of (I) an input layer, (II) a hidden layer consisting of a pattern layer and a competitive layer, and (III) an output layer (see **Figure 3b**, [Demuth2001] and [Specht1990]). The pattern layer contains one unit for each sample in the training dataset. Given an unseen training sample  $x$ , each unit in the pattern layer computes a distance from  $x$  to a specific training instance and applies a Gaussian density activation function. The competitive layer contains one unit for each diagnostic category, and these units receive inputs only from pattern units that are associated with the category to which the training instance belongs. Each of unit in the competitive layer sums over the outputs of the pattern layer and computes a probability of  $x$  belonging to a specific diagnostic category. Finally, the output unit corresponding to maximum of these probabilities outputs 1, while those remaining output 0.

### 2.3. Ensemble classification methods

Given that learners used in this study are different in a sense that they give preference to different models, the final classification performance may be improved via use of algorithms that combine outputs of individual classifiers, so-called *ensembles of classifiers*. This idea has received much attention in machine learning literature (e.g., [Ho1994] and [Sharkey1996]) and has been recently applied to the gene expression domain ([Dudoit2002] and [Valentini2003]). Learning how to combine classifiers to further improve performance is an additional meta-learning problem. Since there is no consensus on which methods are the best in ensembling classifiers, we considered a number of techniques: the most common approach by majority voting [Freund1995] and more complex approaches, Decision Trees (DT) [Murthy1997] and MC-SVM methods (OVR, OVO, DAGSVM). When algorithms were applied for ensembling of classifiers, the input dataset consisted of attributes corresponding to the outputs of classifiers (either SVM or both SVM and non-SVM algorithms) and the original class labels. Combining

classifiers by DT or MC-SVM methods could yield majority voting for some cases, but DT or MC-SVMs allow many more ways to construct ensemble of classifiers.

## 2.4. Parameters for the classification algorithms

Parameters for the classification algorithms were chosen by nested cross-validation procedures to optimize performance while avoiding overfitting as described in the experimental design subsection.

For all five MC-SVM methods we used a polynomial kernel  $K(x, y) = (\gamma \cdot x^T y + r)^p$ , where  $x$  and  $y$  are samples with gene expression values and  $p, \gamma, r$  are kernel parameters. We performed classifier optimization over the set of values of cost  $C$  (the penalty parameter of SVMs) =  $\{0.0001, 0.01, 1, 100\}$  and  $p = \{1, 2, 3\}$ . The kernel parameters  $\gamma$  and  $r$  were set to default values as in [Chang2003]:  $\gamma = 1/\text{number of variables}$  and  $r = 0$ . For Backpropagation Neural Networks, we performed optimization by implementing early stopping regularization techniques following [Goodman1996] on top of the Matlab Neural Network toolbox with parameter selection in a nested cross-validation fashion in order to avoid overfitting. In particular, we used feed-forward NN with one hidden layer and the number of units chosen from the set  $\{2, 5, 10, 30, 50\}$  based on cross-validation error. We employed gradient descent with adaptive learning rate backpropagation, mean squared error performance goal set to  $10^{-8}$  (an arbitrary value very close to zero), fixed momentum of  $10^{-3}$ , and an optimal number of epochs in the range  $[100, 10000]$  based on the early stopping criterion of [Goodman1996]. For Probabilistic Neural Networks, we optimized the smoothing factor  $\sigma$ , a parameter of the Gaussian density function, over 100 different values ranging from 0.01 to 1.00. The parameter  $\sigma$  was set the same for all diagnostic categories. Similarly, we performed a thorough optimization of the KNN classifier over all possible numbers of neighbors  $K$  ranging from 1 to the total number of instances in the training dataset based on cross-validation error.

## 2.5. Datasets and data preparatory steps

The datasets used in this work are described in **Table 1**. In addition to nine multicategory datasets which were most of the multicategory cancer diagnosis datasets in humans found in the public domain at the time the present study was initiated, two binary datasets (i.e. with two diagnoses), *DLBCL* and *Prostate\_Tumor*, were also included to empirically confirm that the employed MC-SVM learners behave well in binary classification tasks as theoretically expected.

The studied datasets were produced primarily by oligonucleotide-based technology. Specifically, in all datasets except for *SRBCT*, RNA was hybridized to high-density oligonucleotide Affymetrix arrays HG-U95 or Hu6800, and expression values (average difference units) were computed using Affymetrix GENECHIP analysis software. The *SRBCT* dataset was obtained by use of two-color cDNA platform with consecutive image analysis performed by DeArray Software and filtering for a minimal level of expression [Khan2001].

The genes or oligonucleotides with “absent” calls in all samples were excluded from the analysis to reduce the amount of noise in the datasets ([Lu2002],[Wouters2003]), and if this was the case, the number of genes is listed in bold in **Table 1**. While setting up datasets for experiments, we took advantage of all available documentation in order to increase the number of categories or diagnoses for the outcome variable. For example, the original *Brain\_Tumor1* data analysis had only two categories – glioblastomas and anaplastic oligodendrogliomas. Instead of a binary classification problem, we solved a diagnostic problem with four outcomes: classic glioblastomas, non-classic glioblastomas, classic anaplastic oligodendrogliomas, and non-classic anaplastic oligodendrogliomas.

**Table 1.** Cancer-related human gene expression datasets used in this study. In addition to 9 multicategory datasets, 2 datasets with two diagnoses were included to empirically confirm that MC-SVM methods behave as well as binary SVMs in binary classification tasks as theoretically expected. The column “Max. prior” indicates the prior probability of the dominant diagnostic category.

Dataset name	Diagnostic Task	Number of				Max. prior	Reference
		Sam- ples	Variables (genes)	Cate- gories	Variables / Samples		
<i>11_Tumors</i>	11 various human tumor types	174	12533	11	72	15.5%	[Su2001]
<i>14_Tumors</i>	14 various human tumor types and 12 normal tissue types	308	<b>15009</b>	26	49	9.7%	[Ramaswamy2001]
<i>9_Tumors</i>	9 various human tumor types	60	<b>5726</b>	9	95	15.0%	[Staunton2001]
<i>Brain_Tumor1</i>	5 human brain tumor types	90	<b>5920</b>	5	66	66.7%	[Pomeroy2002]
<i>Brain_Tumor2</i>	4 malignant glioma types	50	<b>10367</b>	4	207	30.0%	[Nutt2003]
<i>Leukemia1</i>	Acute myelogenous leukemia (AML), acute lymphoblastic leukemia (ALL) B-cell, and ALL T-cell	72	<b>5327</b>	3	74	52.8%	[Golub1999]
<i>Leukemia2</i>	AML, ALL, and mixed-lineage leukemia (MLL)	72	<b>11225</b>	3	156	38.9%	[Armstrong2002]
<i>Lung_Cancer</i>	4 lung cancer types and normal tissues	203	12600	5	62	68.5%	[Bhattacharjee2001]
<i>SRBCT</i>	Small, round blue cell tumors (SRBCT) of childhood	83	2308	4	28	34.9%	[Khan2001]
<i>Prostate_Tumor</i>	Prostate tumor and normal tissues	102	<b>10509</b>	2	103	51.0%	[Singh2002]
<i>DLBCL</i>	Diffuse large b-cell lymphomas (DLBCL) and follicular lymphomas	77	<b>5469</b>	2	71	75.3%	[Shipp2002]

In summary, the 11 datasets had 2-26 distinct diagnostic categories, 50-308 samples (patients), and 2308-15009 variables (genes) after the data preparatory steps outlined above. All datasets are available for download from [Statnikov2004].

We note that no new methods to preprocess gene expression data were invented. We relied instead on standard normalization and data preparatory steps performed by the authors of the primary dataset studies. In addition to that, we performed a simple rescaling of gene expression values to be between 0 and 1 for speeding up SVM training. The rescaling was performed based on the training set in order to avoid overfitting.

## 2.6. Experimental design for model selection and evaluation

Two experimental designs were employed to obtain reliable performance estimates and avoid overfitting. Both experimental designs are based on two loops. The inner loop is used to determine the best parameters of the classifier (i.e. values of parameters yielding the best performance on the validation dataset). The outer loop is used for estimating the performance of the classifier built using the previously found best parameters by testing on an *independent set of patients*. **Design I** uses a stratified 10-fold cross-validation in the outer loop and a stratified 9-fold cross-validation in the inner loop [Weiss1991]. It is often referred to as *nested stratified 10-fold cross-validation*. See **Figure 4** for a simplified pictorial example of a 3-fold Design I applied to 3 patient groups (P1, P2, P3) with optimization of parameter  $C$  (which takes values “1” and “2”) of some classifier. Note, that in reality we do not optimize just one parameter but, rather, a large set of combined parameters. **Design II** uses leave-one-out cross-validation (LOOCV) in the outer loop and a stratified 10-fold cross-validation in the inner loop. We chose to employ both designs because there exists contradictory evidence in the machine learning literature regarding whether  $N$ -fold cross-validation provides more accurate performance estimates than LOOCV and vice-versa for zero-one loss classification [Kohavi1995].

Building of the final diagnostic model involves: (a) finding the best parameters for the classification algorithm using a single loop of cross-validation analogously to the inner loop in Designs I and II; (b) building the classifier on all data using the previously found best parameters; and (c) estimating a conservative bound on the classifier’s future accuracy by running either Design I or II.

## 2.7. Gene selection

To study how dimensionality reduction can improve classification performance, we applied all classifiers with subsets of 25, 50, 100, 500, and 1000 top-ranked genes, following the example set by [Furey2000]. Genes were selected according to four gene selection methods/metrics: (1) ratio of genes between-categories to within-category sums of squares (BW) [Dudoit2002]; (2-3) signal-to-noise (S2N) scores [Golub1999] applied in a one-versus-rest (S2N-OVR) and one-versus-one (S2N-OVO) fashion; and (4) Kruskal-Wallis nonparametric one-way ANOVA (KW) [Jones1997]. The ranking of the genes was performed based on the training set of samples to avoid overfitting.

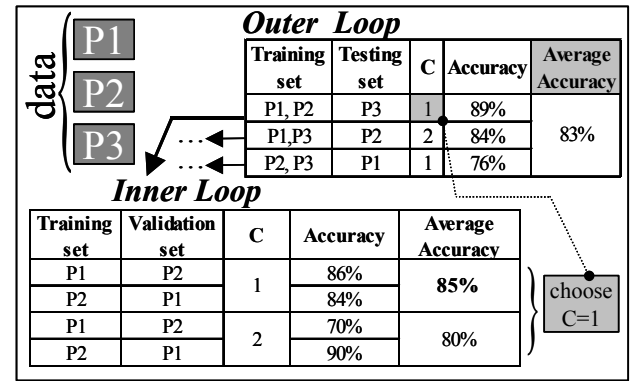
## 2.8. Performance metrics

We used two classification performance metrics. The first metric is accuracy since we wanted to compare our results with the previously published studies that also used this performance metric. Accuracy is easy to interpret and simplifies statistical testing. On the other hand, accuracy is sensitive to the prior class probabilities and does not fully describe the actual difficulty of the decision problem for highly unbalanced distributions. For example, it is more difficult to achieve an accuracy of 50% for a 26-class dataset *14\_Tumors* with prior probability of the major class=9.7% compared to an accuracy of 75% for a binary dataset *DLBCL* with prior of the major class=75.3%.

The second metric is *relative classifier information* (RCI), which corrects for differences in prior probabilities of the diagnostic categories, as well as the number of categories. RCI is an entropy-based measure that quantifies *how much the uncertainty of a decision problem is reduced by a classifier relative to classifying using the priors* [Sindwani2001].

## 2.9. Overall research design

To maintain the feasibility of this study, we pursued a *staged factorial design*: in stage I, we conducted a fully factorial design involving datasets and classifiers without gene selection; in stage II, we focused on the datasets for which the full gene sets yielded poor performance and applied gene selection in a factorial fashion. In addition, we optimized algorithms using accu-



**Figure 4.** Pictorial simplified example of Design I. The data are split into mutually exclusive sets P1, P2 and P3. The performance is estimated in the outer loop by training on all splits but one, using the remaining one for testing. The average performance over testing sets is reported. The inner loop is used to determine the optimal value of parameter  $C$  (in a cross-validated fashion) for training in the outer loop.

racy only and limited the possible cardinalities of selected gene sets to only five choices as described in the subsection on gene selection.

While the above choices restricted the number of models generated, *the resulting analyses still generated more than  $2.6 \cdot 10^6$  diagnostic models*. The total time required was 4 single-CPU months using Intel Xeon 2.4 GHz platform. Out of this set of models, only one model was selected for each combination of algorithm and dataset.

Notice that, despite the very large number of examined models, the final performance estimates are not overfitted. This is because only one model is selected per split for the estimation of the final performance and it is applied to previously *unseen cases*. Thus, regardless of how much performance is overestimated in the inner loop (which, in the worst case, may result in not choosing the best possible parameters' combination), the outer loop guarantees proper estimation of performance.

## 2.10. Statistical comparison among classifiers

To test that differences in accuracy between the best method (i.e. one with the largest average accuracy) and all remaining algorithms are non-random, we need a statistical comparison of observed differences in accuracies.

In machine learning, the major study about comparison of supervised classification algorithms is that of Dietterich which suggests using  $N$ -fold cross-validated paired  $t$ -test for comparison of  $N$ -fold accuracy estimates for a single dataset [Dietterich1998]. However, the author clearly admits that this test violates independence and, even more importantly, does not address how this procedure is applied to a multitude of datasets. That is why we decided to use random permutation testing which does not rely on independence assumptions and can be straightforwardly applied to several datasets [Good2000]. For every algorithm  $X$ , other than the best algorithm  $Y$ , we performed the following steps: (I) We defined the null hypothesis  $H_0$  to be: classification algorithm  $X$  is as good as  $Y$ , i.e. the accuracy of the best algorithm  $Y$  minus the accuracy of algorithm  $X$  is zero. (II) We obtained the permutation distribution of  $\Delta_{XY}$ , the estimator of the true unknown difference between accuracies of the two algorithms, by repeatedly rearranging the outcomes of  $X$  and  $Y$  at random. (III) We computed the cumulative probability (p-value) of  $\Delta_{XY}$  being greater than or equal to observed difference  $\hat{\Delta}_{XY}$  over 10,000 permutations. If the p-value was smaller than 0.05, we rejected  $H_0$  and concluded that the data support that algorithm  $X$  is not as good as  $Y$  in terms of classification accuracy, and this difference is not due to sampling error. In order to increase the resolution of simulated sampling distribution, we computed a single value of accuracy over all samples from all datasets. In other words, we treated classifier's predictions from all 11 datasets as if we had one large dataset with samples from all individual datasets.

## 3. Implementations

We used the MC-SVM algorithms implemented by the LibSVM team [Chang2003], since they use state-of-the-art optimization methods SMO [Platt1999] and TRON [Lin1999] for the solution of MC-SVM problems. The implementation of NN and PNN classifiers was based on the Matlab Neural Networks toolbox [Demuth2001]. We applied Matlab R13 implementation of the CART algorithm [Murthy1997] for DT, and we used our own implementations of KNN, ensemble classification, gene selection, as well as statistical comparison algorithms.

The prototype analysis system **GEMS** (Gene Expression Model Selector) based on the results and analyses reported here, was built using Matlab R13 and MS Visual C++ 6. **GEMS** has a graphics user interface consisting of a single form (Figure 5) and is freely available for download from [Statnikov2004]. The user's manual for the system is provided in Appendix D, which is available online from [Statnikov2004].

**Figure 5.** Screenshot of the GEMS system. Many fields are automatically filled out with default values. Most experiments in this study can be replicated using the system with a few clicks of the mouse.

## 4. Results and Analyses

### 4.1. Classification without gene selection

The performance results of experiments without gene selection obtained using Design I (nested stratified 10-fold cross-validation) with accuracy and RCI as a performance metric are shown in **Tables 2** and **3**, respectively. Results for Design II are almost identical and are provided only in Appendix E, section 1, which is available online from [Statnikov2004]. The fact that we obtained similar results with two different experimental designs is evidence in favor of the reliability of performance estimation procedures.

Notably, RCI performance metric revealed different results compared to accuracy. For example, the best RCI for *14\_Tumors* dataset is 90.96% and for *Prostate\_Tumor* is 71.14%. In contrast, when accuracy was employed, we obtained 76.60% in *14\_Tumors* and 92% in *Prostate\_Tumor*. The difference can be explained by the difficulties of the classification problems - *14\_Tumors* is much harder (it has 26 classes with prior of the most frequent class 9.7%, see **Table 1**) than *Prostate\_Tumor* (it is a binary problem with prior 51%, see **Table 1**).

According to **Table 2**, in 8 out of 11 datasets, MC-SVMs perform cancer diagnoses with accuracies > 90%. The results for RCI performance metric are similar (**Table 3**): in 7 out of 11 datasets, MC-SVMs yield diagnostic performance with RCI > 90%. Overall, all MC-SVMs outperform KNN, NN, and PNN significantly. The only exception is KNN and PNN applied to *14\_Tumors* dataset

which outperformed OVO and DAGSVM, but still were unable to perform better than more robust MC-SVM techniques, OVR, WW, and CS. The superior classification performance of the SVM-based methods compared to KNN, NN, and PNN reflects that SVMs are less sensitive to the *curse of dimensionality* and more robust to a small number of high-dimensional gene expression samples than other non-SVM techniques [Aliferis2003b]. A more detailed explanation of this matter follows in the next subsection on classification results with gene selection.

**Table 2.** Performance results (accuracies) without gene selection obtained using a nested stratified 10-fold cross-validation design (Design I). These results are further improved by gene selection (see **Figure 6**). The last column in the bottom table reports average performance computed over datasets.

Multicategory classification						
Method	9 Tumors	11 Tumors	14 Tumors	Brain Tumor1	Brain Tumor2	Leukemia1
MC-SVM	OVR	65.10%	94.68%	74.98%	<b>91.67%</b>	77.00%
	OVO	58.57%	90.36%	47.07%	90.56%	<b>77.83%</b>
	DAGSVM	60.24%	90.36%	47.35%	90.56%	<b>77.83%</b>
	WW	62.24%	94.68%	69.07%	90.56%	73.33%
	CS	<b>65.33%</b>	<b>95.30%</b>	<b>76.60%</b>	90.56%	72.83%
non-SVM	KNN	43.90%	78.51%	50.40%	87.94%	68.67%
	NN	19.38%	54.14%	11.12%	84.72%	60.33%
	PNN	34.00%	77.21%	49.09%	79.61%	62.83%

Multicategory classification						Averages
Method	Leukemia2	Lung Cancer	SRBCT	Prostate Tumor	DLBCL	
MC-SVM	OVR	<b>97.32%</b>	96.05%	<b>100.00%</b>	<b>92.00%</b>	<b>97.50%</b>
	OVO	95.89%	95.59%	<b>100.00%</b>	<b>92.00%</b>	<b>97.50%</b>
	DAGSVM	95.89%	95.59%	<b>100.00%</b>	<b>92.00%</b>	<b>97.50%</b>
	WW	95.89%	95.55%	<b>100.00%</b>	<b>92.00%</b>	<b>97.50%</b>
	CS	95.89%	<b>96.55%</b>	<b>100.00%</b>	<b>92.00%</b>	<b>97.50%</b>
non-SVM	KNN	87.14%	89.64%	86.90%	85.09%	86.96%
	NN	91.03%	87.80%	91.03%	79.18%	89.64%
	PNN	83.21%	85.66%	79.50%	79.18%	80.89%

**Table 3.** Performance results (RCI) without gene selection obtained using a nested stratified 10-fold cross-validation design (Design I). These results are further improved by gene selection (see **Figure 7**). The last column in the bottom table reports average performance computed over datasets.

Multicategory classification						
Method	9 Tumors	11 Tumors	14 Tumors	Brain Tumor1	Brain Tumor2	Leukemia1
MC-SVM	OVR	77.00%	95.80%	90.53%	<b>82.31%</b>	77.49%
	OVO	78.24%	92.24%	64.99%	80.77%	<b>80.27%</b>
	DAGSVM	<b>78.67%</b>	92.24%	65.64%	80.77%	<b>80.27%</b>
	WW	76.22%	95.80%	86.30%	80.77%	74.75%
	CS	77.25%	<b>96.20%</b>	<b>90.96%</b>	80.77%	74.44%
non-SVM	KNN	63.38%	83.93%	82.73%	67.86%	64.48%
	NN	65.57%	67.80%	16.24%	61.42%	62.49%
	PNN	55.59%	81.39%	81.40%	43.86%	61.73%

Multicategory classification						Averages
Method	Leukemia2	Lung Cancer	SRBCT	Prostate Tumor	DLBCL	
MC-SVM	OVR	<b>94.42%</b>	89.45%	<b>100.00%</b>	<b>71.14%</b>	<b>90.91%</b>
	OVO	92.35%	87.95%	<b>100.00%</b>	<b>71.14%</b>	<b>90.91%</b>
	DAGSVM	92.35%	87.95%	<b>100.00%</b>	<b>71.14%</b>	<b>90.91%</b>
	WW	91.90%	87.46%	<b>100.00%</b>	<b>71.14%</b>	<b>90.91%</b>
	CS	91.90%	<b>91.40%</b>	<b>100.00%</b>	<b>71.14%</b>	<b>90.91%</b>
non-SVM	KNN	76.95%	68.48%	80.71%	51.09%	63.08%
	NN	78.02%	64.97%	87.50%	33.25%	58.36%
	PNN	73.51%	59.72%	68.92%	39.22%	38.23%

Among MC-SVMs, OVR, WW, and CS yield the best results and are not statistically significant from each other at the 0.05 level (Appendix E, section 2). On the other hand, OVO, DAGSVM, KNN, PNN, and NN have poorer performance than the above methods to a statistically significant degree. OVO and DAGSVM perform very similar, which is due to the fact that both MC-SVM methods use the same binary SVM classifiers. We conjecture that OVO and DAGSVM perform worse than other MC-SVM methods because both algorithms are based on one-versus-one binary classifiers that use only a fraction of total training samples at a time (samples that belong to two classes) and ignore information about distribution of the remaining examples which may be significant for the classification. In case of large sample sizes, we expect MC-SVMs OVO and DAGSVM to perform as good as WW, CS, and OVR (for example, see [Hsu2002]).

According to **Tables 2 and 3** and results of application of the binary SVM implementation SVMlight [Joachims1999] to *DLBCL* and *Prostate\_Tumor* datasets (not shown here), we conclude that our implementations of MC-SVM algorithms perform the same classifications as binary SVMs and, hence, handle binary diagnostic problems appropriately as expected.

We tried to explain classification performance of the best MC-SVM algorithms OVR, WW, and CS by fitting inverse power curves motivated by the ideas described in [Cortes1993]. We found that in high-dimensional spaces of microarray gene expression data, the number of samples divided by the product of the number of variables times the number of categories explains observed classification accuracies in the datasets. When we reduced dimensionality by gene selection, or employed RCI performance metric, or used other classification algorithms, this behavior disappeared. More details can be found in Appendix E, section 3. It is important to note that curve fitting procedure used in this study is very simplistic since it does not incorporate predictors describing degree of biological difficulty and assumes that datasets and learning tasks used in this study are representative.

Finally, we also analyzed execution time for all learning algorithms applied without gene selection (**Table 4**). The fastest MC-SVM methods CS and WW took 7.95 and 7.88 hours for Design I and 289.01 and 290.77 hours for Design II, respectively. The slowest MC-SVM technique OVR completed within 19.28 hours for Design I and 772.43 hours for Design II. This technique is slowest among MC-SVM algorithms since it constructs several classifiers repeatedly employing all samples from the training dataset. The fastest overall algorithm KNN took 3.40 hours for Design I and 109.60 hours for Design II, while the slowest overall algorithms NN and PNN took 195.68 hours and 186.19 hours, respectively, for Design I. All experiments were executed in the Matlab R13 environment on eight Intel Xeon 2.4GHz dual-CPU workstations connected in a cluster.

#### 4.2. Classification with gene selection

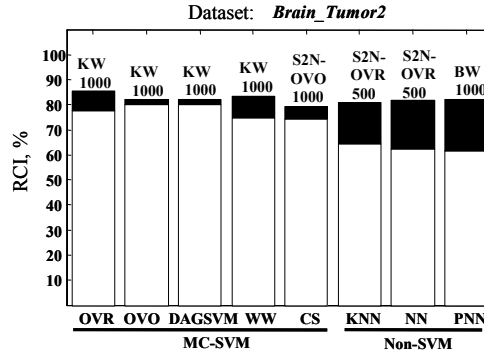
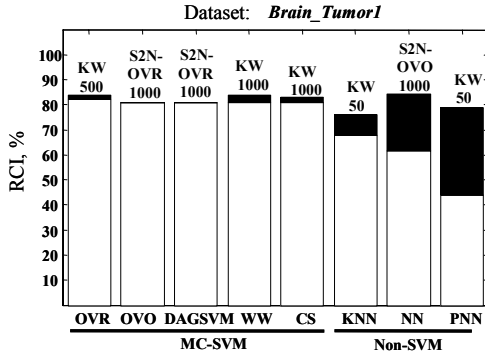
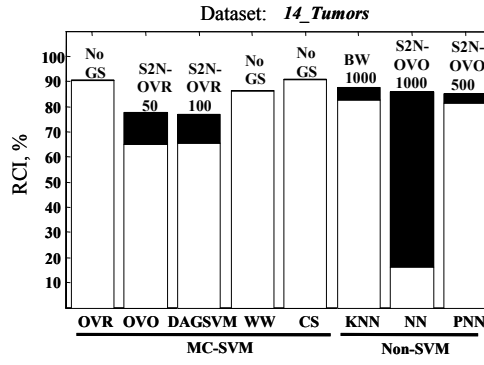
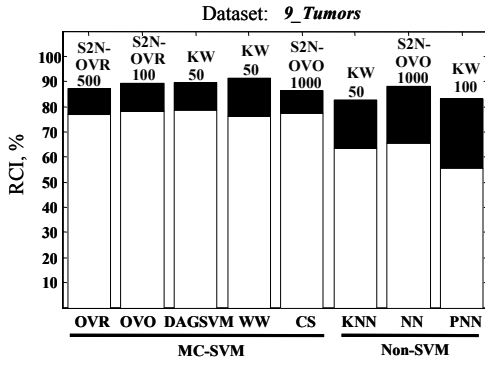
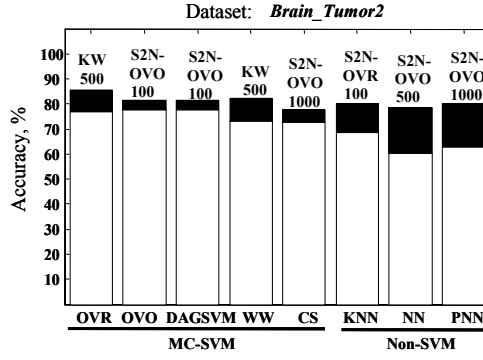
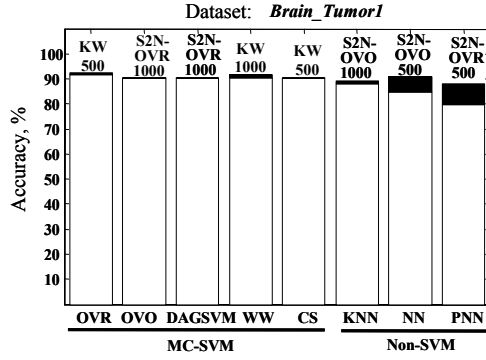
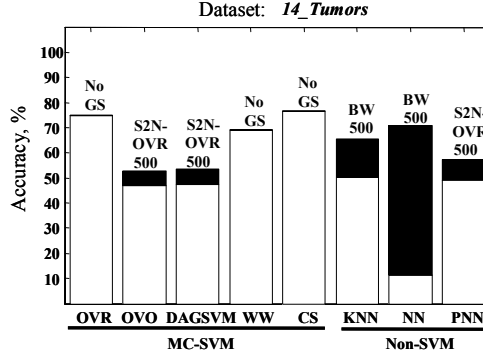
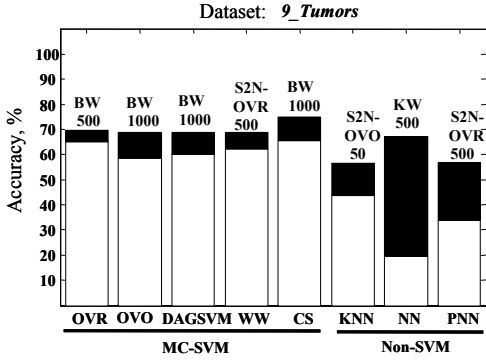
The summary of application of the four gene selection methods BW, S2N-OVR, S2N-OVO, and KW to the four most “challenging” datasets *9\_Tumors*, *14\_Tumors*, *Brain\_Tumor1*, and *Brain\_Tumor2* using accuracy and RCI as a performance metric is presented in **Figures 6 and 7**, respectively. It should be noted that a more rigorous way to do gene selection with validation of number of genes and gene selection method is implemented in *GEMS* software system and may be very expensive computationally (that is why it was not pursued here as explained in the methods section).

The results show that gene selection significantly improves classification performance of non-SVM learners. In particular, for some datasets, accuracy is improved by up to 14.97%, 59.78%, 22.67% and RCI is improved by up to 19.52%, 69.95%, 34.98% for KNN, NN, and PNN, respectively. Gene selection also improves accuracy of MC-SVMs up to 9.53% and, hence, improves accuracy of the overall best classifier. Although KNN, NN, and PNN perform closer to MC-SVMs, three MC-SVM algorithms, OVR, WW, and CS, still outperform non-SVM methods in most of cases. We also found that these three MC-SVM methods are not statistically significant from each other and NN at the 0.05 level (Appendix E, section 2). The remaining algorithms, MC-SVMs OVO and DAGSVM, KNN and PNN, have statistically significant poorer performance. Finally, neither of the four gene selection methods performs significantly better than the other ones.

As we have empirically found, the non-SVM methods KNN, PNN, and NN benefit significantly more than MC-SVMs from gene selection. A number of observations can explain this behavior: In high-dimensional spaces, KNN has high variance of the prediction since all training points are located close to the edge of the sample [Hastie2001]. Furthermore, many irrelevant variables in the data dominate distances between samples which presents a significant problem for prediction [Mitchell1997]. PNN encounter problems similar to KNN, in particular because they rely on Parzen windows for density estimation which generally require exponential sample to the data dimensionality [Duda2001]. Backpropagation Neural Networks are sensitive to high dimensionality for at least two reasons: First notice, that the larger the number of variables, the larger is the number of weights in this type of neural network. Because of this, (I) there may be more local minima in the error landscape and it is thus more probable for backpropagation to get “trapped” in one of them, and (II) the model space becomes exponentially larger with the addition of each weight, and so it becomes harder to identify a model that generalizes. In comparison, the family of SVMs allows for effective optimization search procedure by utilizing convex formulation with a single optimum justified by Statistical Learning Theory [Vapnik1998]. Furthermore, SVMs seem relatively insensitive to the curse of dimensionality, possibly due to the specific regularization mechanism they employ. In particular, this is reflected by the following: (I) many es-

**Table 4.** Total time of classification experiments without gene selection for all 11 datasets and two experimental designs.

Method	Time in hours	
	Design I	Design II
MC-SVM	OVR	19.28
	OVO	9.86
	DAGSVM	9.93
	WW	7.95
	CS	7.88
non-SVM	KNN	3.40
	NN	195.68
	PNN	186.19



**Figure 6.** Performance results (accuracies) of classification experiments with gene selection obtained using a nested stratified 10-fold cross-validation design (Design I) for 4 datasets: 9\_Tumors, 14\_Tumors, Brain\_Tumor1, Brain\_Tumor2. The white bars correspond to classification results without gene selection. The black bars correspond to results with gene selection. The text above each bar indicates the optimal combination of gene selection method and number of genes for a specific classifier. The abbreviation “No GS” stands for “No gene selection”.

**Figure 7.** Performance results (RCI) of classification experiments with gene selection obtained using a nested stratified 10-fold cross-validation design (Design I) for 4 datasets: 9\_Tumors, 14\_Tumors, Brain\_Tumor1, Brain\_Tumor2. The white bars correspond to classification results without gene selection. The black bars correspond to results with gene selection. The text above each bar indicates the optimal combination of gene selection method and number of genes for a specific classifier. The abbreviation “No GS” stands for “No gene selection”.

established generalization bounds do not depend on the data dimensionality [Herbrich2002], and (II) even linear SVMs assign zero weights to irrelevant variables [Hardin2004]. On the other hand, the SVM algorithm may assign non-zero weights to weakly relevant variables [Hardin2004] which explains why effective variable selection can still improve SVM classification.

### 4.3. Ensemble classification

For the case when no gene selection was performed, ensembles do not outperform the best non-ensemble methods with the exception of the Decision Trees ensemble classifier for *Brain\_Tumor2* dataset, which improves classification accuracy by 1.67%. Other ensembles often achieve similar performance to the best non-ensemble methods (Appendix E, section 4).

Next, we considered three datasets 9\_Tumors, *Brain\_Tumor1*, and *Brain\_Tumor2* where we previously observed improvement of classification performance by gene selection. For each dataset we selected a subset of genes yielding the best classification performance (over all gene selection methods, subsets of genes, and learning algorithms) and constructed combined classifiers. According to results, ensembles perform worse than the best non-ensemble models (Appendix E, section 4).

We believe that in our study, ensemble classifiers did not improve final classification performance for the following two reasons: First, samples misclassified by non-SVM algorithms are almost always a strict superset of samples misclassified by MC-SVM algorithms. Second, SVM algorithms are fairly stable in a sense that small changes in the training data do not result in large changes in the predictive model’s behavior [Kutin2002], and according to [Dudoit2002] stable algorithms do not usually tend to benefit from the ensemble classification.

#### 4.4. Comparison with previously published results

Most of the results from this study are not exactly comparable with the analyses provided in the original studies due to differences in the setup of dataset/learning task, experimental design, gene selection, classifiers, etc. that vary from study to study. However, the reported results in the literature confirm that MC-SVMs as applied here perform equally as well, or even better, compared to previously published models on the same datasets (Appendix E, section 5).

### 5. Discussion and Limitations

One of the limitations of the present study is that we use accuracy and RCI as our performance measures. These metrics do not incorporate information about confidence of the predictions as well as different misclassification costs of diagnostic categories. On the other hand, accuracy was used in published studies and it is easy to interpret and simplifies statistical comparison, while RCI is insensitive to prior class probabilities and accounts for the difficulty of the learning problem. There are currently no mature performance metrics applicable for multiclass domains and suitable for our classifiers with both confidence information and consideration of misclassification costs. Initial attempts were introduced by [Lee2003], [Mossman1999], and [Ferri2003], however much needs to be done before we obtain a workable metric for experiments such as those presented here.

As we mentioned, the choice of KNN, NN, and PNN classifiers as the baseline techniques was grounded on prior successful applications to gene expression based cancer diagnosis (e.g., [Khan2001], [Ramaswamy2001], [Pomeroy2002], [Nutt2003], [Singh2002], and [Berrar2003]). We have also experimented with other non-SVM classifiers, such as Decision Trees (DT) [Murthy1997] and Weighed Voting (WV) classifiers applied both in OVR and OVO fashion ([Golub1999], [Ramaswamy2001], and [Yeang2001]). We found that both with and without gene selection, DT perform significantly worse than MC-SVMs, worse than KNN, and similarly or worse than NN and PNN. Likewise, WV classifiers are significantly outperformed by MC-SVMs, KNN, NN, and PNN. More details about these additional experiments with DT and WV classifiers can be found in Appendix E, section 6.

A particularly interesting direction for future research is to improve our existing gene selection procedures with selection of the “optimal” number of genes by cross-validation<sup>1</sup>. Furthermore, we are interested in applying various multivariate Markov blanket and local neighborhood algorithms which have been previously successfully applied to cancer gene expression and several other domains and do guarantee efficient identification of a set of relevant attributes under fairly broad assumptions ([Aliferis2003c], [Tsamardinos2003]).

We plan to extend our comparative analyses with new MC-SVM methods as they become available. In particular, we plan to use MSVM<sup>2</sup> (see [Lee2003] for details) which has promising theoretical properties. Namely, MSVM employs a certain loss function for the multicategory classification problem under which the solution to the multicategory problem resembles Bayes rule asymptotically. Moreover, this framework easily allows accounting for unequal misclassification and distortion of class proportions. In addition to new MC-SVM algorithms, a promising set of new methods is error correcting output codes for solution of multiclass problems by reducing them to binary problems<sup>3</sup> that can be solved using binary SVMs. This approach is very promising in that some researchers proved a general empirical multiclass error loss bound given empirical loss of individual binary classifiers [Allwein2000].

The emergence of new cancer gene expression datasets in our institution and elsewhere will allow us to conduct a prospective evaluation of the *GEMS* system to study its ability to facilitate creation of powerful diagnosis models. We are also working on augmenting the preliminary version of the system with wizard-like graphics user interface that will make *GEMS* usable by researchers with limited expertise in data-analysis.

To the best of our knowledge, currently there exists only one work aimed at evaluation of MC-SVM algorithms [Hsu2002]. That study is outside of the realm of biomedicine since [Hsu2002] considered such classification tasks as wine recognition, letter recognition, shuttle control, etc. with the number of variables ranging from 4 to 180 and sample sizes greater than 500 in the majority of tasks, which is not typical for microarray cancer gene expression datasets. However, it is worthwhile to mention the major conclusions of that evaluation. The authors empirically found the following: (1) using a Gaussian radial basis

<sup>1</sup> The functionality to cross-validate number of genes is already implemented in the software system.

<sup>2</sup> Although the codes of MSVM were kindly provided by its author, the implementation of this algorithm cannot currently handle problems with very large number of categories and/or large sample size. Given the excellent theoretical properties of the MSVM, we hope that this issue will be solved in the near future so this algorithm can be applied to the present problem domain.

<sup>3</sup> In the present study we have already employed three approaches to reduce multiclass problems to binary – OVR, OVO, and DAGSVM.

kernel, all MC-SVM methods perform similarly; (2) DAGSVM and OVO have the fastest training time; and (3) for problems with large sample size, WW and CS yield fewer support vectors compared to OVR, OVO, and DAGSVM. The work by Hsu is complementary to ours and is not overlapping due to significant differences in the problem domain and dataset characteristics. For example, in our experiments, MC-SVM methods OVO and DAGSVM achieved inferior classification performance compared to other MC-SVM algorithms.

Finally, two recent bioinformatics studies have also performed comparative analyses of multicategory classification algorithms in cancer gene expression domain ([Berrar2003] and [Romualdi2003]). Unfortunately, neither study optimized parameters of the classifiers for all datasets, which is likely to result in suboptimal application of the learning methods. Therefore, neither study can convincingly answer the central question of this research - what is the best learning algorithm for multicategory cancer diagnosis based on gene expression data?

## 6. Conclusions

The contributions of the present study are two-fold. The *first contribution* is that we conducted the most comprehensive systematic evaluation to date of multicategory diagnosis algorithms applied to the majority of multicategory cancer-related gene expression human datasets publicly available. Based on results of this evaluation, the following conclusions can be drawn:

- *Multicategory Support Vector Machines is the best family of algorithms for this type of data and medical tasks.* They outperform other popular non-SVM machine learning techniques by a large margin.
- Among MC-SVM methods, the ones by Crammer and Singer, Weston and Watkins, and one-versus-rest have superior classification performance.
- The performance of both MC-SVM and non-SVM methods can be moderately (for MC-SVMs) or significantly (for non-SVM) improved by gene selection.
- Ensemble classification does not further improve classification performance of the best MC-SVM models.

We believe that practitioners and software developers should take note of these results when considering construction of decision support systems in this domain, or when selecting algorithms for inclusion in related analysis software.

The *second contribution* is that we created the fully-automated software system *GEMS* that automates the experimental procedures described in this paper to (1) develop optimal classification models for the domain of cancer diagnosis with microarray gene expression data and (2) estimate their performance in future patients. The results obtained by the system in a labor-efficient manner appear to be on par with or better than previously published results in the literature on the same datasets. Although several commercial and academic software tools exist for gene expression classification (e.g., [Reich2004]), *to the best of our knowledge GEMS treats the task in the most comprehensive manner and is the first such system to be informed by a rigorous analysis of the available algorithms and datasets.* We hope that the methodology presented in the present paper may encourage similar principled treatment of other software development efforts in clinical bioinformatics. The system is freely available for download from [Statnikov2004] for non-commercial use and its user's manual is provided in Appendix D.

## 7. Acknowledgements

This research was supported by NIH grants RO1 LM007948-01 and P20 LM 007613-01.

## References

- [Aliferis2003a] Aliferis C.F., I. Tsamardinos, P. Massion, A. Statnikov, N. Fananapazir, D. Hardin. "Machine Learning Models For Classification Of Lung Cancer and Selection of Genomic Markers Using Array Gene Expression Data", In Proceedings of the 16th International Florida Artificial Intelligence Research Society (FLAIRS) Conference, 2003.
- [Aliferis2003b] Aliferis C.F., I. Tsamardinos, P. Massion, A. Statnikov, D. Hardin. "Why Classification Models Using Array Gene Expression Data Perform So Well: A Preliminary Investigation Of Explanatory Factors", In Proceedings of the 2003 International Conference on Mathematics and Engineering Techniques in Medicine and Biological Sciences (METMBS), 2003.
- [Aliferis2003c] Aliferis C.F., I. Tsamardinos, A. Statnikov. "HITON, A Novel Markov Blanket Algorithm for Optimal Variable Selection", In Proceedings of the 2003 American Medical Informatics Association (AMIA) Annual Symposium, 2003.
- [Allwein2000] Allwein E.L., R. E. Schapire, and Y. Singer. "Reducing Multiclass to Binary: A Unifying Approach for Margin Classifiers", *Journal of Machine Learning Research* 1: 113-141, 2000.
- [Armstrong2002] Armstrong S., et al. "MLL translocations specify a distinct gene expression profile that distinguishes a unique leukemia", *Nature Genetics*, volume 30, January 2002.
- [Berrar2003] Berrar D. "Multiclass Cancer Classification Using Gene Expression Profiling And Probabilistic Neural Networks", In Proceedings of the Pacific Symposium on Biocomputing (PSB), 2003.
- [Bhattacharjee2001] Bhattacharjee, A., et al. "Classification of human lung carcinomas by mRNA expression profiling reveals distinct adenocarcinoma subclasses", *Proc Natl Acad Sci U S A*, 2001 Nov 20;98(24):13790-5.
- [Chang2003] Chang, Chih-Chung and Lin, Chih-Jen. "LIBSVM: a library for Support Vector Machines", <http://www.csie.ntu.edu.tw/~cjlin/libsvm>, 2003.
- [Cortes1993] Cortes, C., Lawrence D. Jackel, Sara A. Solla, Vladimir Vapnik, John S. Denker. "Learning Curves: Asymptotic Values and Rate of Convergence", *Advances in Neural Information Processing Systems (NIPS)* 1993: 327-334.

- [Crammer2000] Crammer, K. and Y. Singer. "On the Learnability and Design of Output Codes for Multiclass Problems", Proceedings of the Thirteen Annual Conference on Computational Learning Theory (COLT), 2000.
- [Demuth2001] Demuth, H. and M. Beale. "Neural network toolbox user's guide. Matlab user's guide", 2001: The MathWorks Inc.
- [Dietterich1998] Dietterich, T. G. "Approximate statistical tests for comparing supervised classification learning algorithms", Neural Computation, 10 (7), 1998.
- [Duda2001] Duda, R.O, Hart, P.E, and Stork, D.G. "Pattern Classification", second edition, John Wiley, New York, 2001.
- [Dudoit2002] Dudoit S., 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, Vol. 97, No. 457, p. 77-87, 2002.
- [Ferri2003] Ferri C., J. Hernández-Orallo, and M.A. Salido. "Volume under the ROC Surface for Multi-Class Problems", In Proc. of 14th European Conference on Machine Learning, ECML'03, LNAI Springer Verlag, Vol 2837, pages 108-120, 2003.
- [Fortina2002] Fortina P, et al. "Molecular diagnostics: hurdles for clinical implementation", Trends Mol Med. 2002 Jun; 8(6): 264-6.
- [Freund1995] Freund, Y. "Boosting a weak learning algorithm by majority", Information and Computation, 121(2):256-285, 1995.
- [Friedman1996] Friedman, J. "Another approach to polychotomous classification", Technical report, Stanford Univeristy, 1996.
- [Furey2000] Furey, T.S., et al.. "Support vector machine classification and validation of cancer tissue samples using microarray expression data", Bioinformatics 2000 16: 906-914.
- [Golub1999] Golub, T., et al. "Molecular Classification of Cancer: Class Discovery and Class Prediction by Gene Expression Monitoring", Science, Vol 286, 15 October 1999.
- [Good2000] Good, P. I. "Permutation Tests: A Practical Guide to Resampling Methods for Testing Hypotheses", 2nd ed. New York: Springer-Verlag, 2000.
- [Goodman1996] Goodman P. and Frank Harrell, "NevProp Manual with Introduction to Artificial Neural Networks Theory" <http://brain.cs.unr.edu/publications/NevPropManual.pdf>, 2004.
- [Guyon2002] Guyon, I., et al. "Gene selection for cancer classification using support vector machines", Machine Learning, 2002, 46: 389-422.
- [Guyon2003] Guyon, I., et al. "Erratum: Gene selection for cancer classification using support vector machines", <http://www.clopinet.com/isabelle/Papers/RFE-erratum.html>, 2003.
- [Hardin2004] Hardin D., I. Tsamardinos, C.F. Aliferis, "A Theoretical Characterization of Linear SVM-Based Feature Selection", In the Twenty-First International Conference on Machine Learning (ICML), 2004.
- [Hastie2001] Hastie T., Robert Tibshirani and Jerome Friedman, "Elements of Statistical Learning: Data Mining, Inference and Prediction" Springer-Verlag, New York, 2001.
- [Herbrich2002] Herbrich, R. "Learning Kernel Classifiers: Theory and Algorithms". MIT Press. 2002.
- [Ho1994] Ho TK, Hull JJ, Srihari SN. "Decision combination in multiple classifier systems", IEEE Trans Pattern Analysis and Machine Intelligence 1994; 16(1):66-76.
- [Hsu2002] Hsu, Chih-Wei and Chih-Jen Lin. "A Comparison of Methods for Multi-class Support Vector Machines", IEEE Transactions in Neural Networks 13(2) 415-425, 2002.
- [Joachims1999] Joachims, T. "Making Large-Scale SVM Learning Practical", Advances in Kernel Methods - Support Vector Learning, B. Schölkopf, C. Burges, and A. Smola (ed.), MIT Press, 1999.
- [Jones1997] Jones, B. "Matlab Statistics Toolbox", The MathWorks, Inc. Natick, MA, USA, 1997.
- [Khan2001] Khan, J., et al. "Classification and diagnostic prediction of cancers using gene expression profiling and artificial neural networks", Nature Medicine, volume 7, Number 6, June 2001.
- [Kohavi1995] Kohavi, R. "A Study of Cross-Validation and Bootstrap for Accuracy Estimation and Model Selection", In Proceedings of the International Joint Conferences on Artificial Intelligence (IJCAI), 1995.
- [Kressel1999] Kressel, U. "Pairwise classification and support vector machines", In Advances in Kernel Methods: Support Vector Learning (Chapter 15), MIT Press, 1999.
- [Kutin2002] Samuel Kutin, Partha Niyogi: Almost-everywhere Algorithmic Stability and Generalization Error. UAI 2002: 275-282.
- [Lee2003] Lee, Yoonkyung and Cheol-Koo Lee. "Classification of multiple cancer types by multcategory support vector machines using gene expression data", Bioinformatics 2003 19: 1132-1139.
- [Lin1999] Lin, Chih-Jen and Jorge J. Moré. "Newton's Method for Large Bound-Constrained Optimization Problems", SIAM Journal on Optimization, Volume 9, Number 4, pp. 1100-1127, 1999.
- [Lu2002] Lu, J., et al. "Classical statistical approaches to molecular classification of cancer from gene expression profiling", in Methods of Microarray Data Analysis: Papers from CAMDA'00, eds. S.M. Lin, K.F. Johnson, pp. 97-107, Kluwer Academic Publishers, Boston/Dordrecht/London, 2002.
- [Mitchell1997] Mitchell, T.M., "Machine Learning", McGraw-Hill, 1997.
- [Mossman1999] Mossman, D. "Three-way ROCs", Medical Decision Making, 1999, 19: 78-89.
- [Mukherjee2003] Mukherjee S. "Classifying Microarray Data Using Support Vector Machines", Understanding And Using Microarray Analysis Techniques: A Practical Guide. Boston: Kluwer Academic Publishers; 2003.
- [Murthy1997] Murthy, S., "Automatic construction of decision trees from data: A multi-disciplinary survey", Data Mining and Knowledge Discovery, 1997.
- [Ntzani2003] Ntzani EE and Ioannidis JP. "Predictive ability of DNA microarrays for cancer outcomes and correlates: and empirical assessment", Lancet. 2003 Nov 1;362(9394):1439-44.
- [Nutt2003] Nutt, C., et al. "Gene expression-based classification of malignant gliomas correlates better with survival than histological classification", Cancer Res. 2003 Apr 1;63(7):1602-7.
- [Platt1999] Platt, J. "Fast Training of Support Vector Machines using Sequential Minimal Optimization", Advances in Kernel Methods - Support Vector Learning, B. Schölkopf, C. Burges, and A. Smola, (ed.), MIT Press, 1999.
- [Platt2000] Platt, J., N. Cristianini, and J. Shawe-Taylor. "Large margin dags for multiclass classification", Advances in Neural Information Processing Systems 12, pages 547-553. MIT Press, 2000.

- [Pomeroy2002] Pomeroy, L., et al. "Prediction of central nervous system embryonal tumour outcome based on gene expression", *Nature*, vol 415, 24 January 2002.
- [Ramawamy2001] Ramaswamy, S., et al. "Multiclass cancer diagnosis using tumor gene expression signatures", *Proc Natl Acad Sci U S A* Dec 11, 2001.
- [Reich2004] Reich M, Ohm K, Angelo M, Tamayo P, Mesirov JP. "GeneCluster 2.0: an advanced toolset for bioarray analysis", *Bioinformatics*. 2004 Jul 22;20(11):1797-8.
- [Reunanen2003] Reunanen, J. "Overfitting in Making Comparisons Between Variable Selection Methods", *Journal of Machine Learning Research* 3 (2003) 1371-1382.
- [Romualdi2003] Romualdi C., Campanaro S., Campagna D., Celegato B., Cannata N., Toppo S., Valle G. and Lanfranchi G. (2003) Pattern recognition in gene expression profiling using DNA array: a comparative study of different statistical methods applied to cancer classification. *Hum. Mol. Gen.* 12:823-36.
- [Schwarzer2000] Schwarzer G and Vach W. "On the misuses of artificial neural networks for prognostic and diagnostic classification in oncology", *Stat Med.* 2000 Feb 29;19(4):541-61.
- [Sindwani2001] Sindwani V., et al. "Information Theoretic Feature Crediting in Multiclass Support Vector Machines", *First SIAM International Conference on Data Mining*, 2001.
- [Singh2002] Singh, D., et al. "Gene expression correlates of clinical prostate cancer behavior", *Cancer Cell* : March 2002, Vol. 1.
- [Sharkey1996] Sharkey AJC (ed). "Special Issue: Combining Artificial Neural Networks: Ensemble Approaches", *Connection Science* 1996; 8(3 & 4).
- [Shipp2002] Shipp, M., et al. "Diffuse large B-cell lymphoma outcome prediction by gene expression profiling and supervised machine learning", *Nature Medicine*, Volume 8, Number 1, January 2002.
- [Specht1990] Specht, DF., "Probabilistic neural network," *Neural Networks*, vol. 3, pp. 109-118, 1990.
- [Statnikov2004] Statnikov A, Aliferis C, Tsamardinos I, Hardin D, and Levy S. <http://www.gems-system.org>, 2004.
- [Staunton2001] Staunton, J., et al. "Chemosensitivity prediction by transcriptional profiling", *Proc Natl Acad Sci U S A*, September 11, 2001, vol. 98, no. 19, 10787-10792.
- [Su2001] Su, A.I., et al. "Molecular classification of human carcinomas by use of gene expression signatures", *Cancer Res.* 2001 Oct 15;61(20):7388-93.
- [Tsamardinos2003] Tsamardinos, I., C.F. Aliferis, A. Statnikov. "Time and Sample Efficient Discovery of Markov Blankets and Direct Causal Relations", 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2003.
- [Valentini2003] Valentini G., M. Muselli, and F. Ruffino. "Bagged Ensembles of SVMs for Gene Expression Data Analysis", *IJCNN2003*, The IEEE-INNS-ENNS International Joint Conference on Neural Networks, Portland, USA, 2003.
- [Vapnik1998] Vapnik, V. "Statistical Learning Theory", Wiley-Interscience, 1998.
- [Weiss1991] Weiss S.M. and C.A. Kulikowski. "Computer systems that learn", Morgan Kaufmann, 1991.
- [Weston1999] Weston, J. and C. Watkins. "Support Vector Machines for Multi-Class Pattern Recognition", *Proceedings of the Seventh European Symposium On Artificial Neural Networks*, 1999.
- [Wouters2003] Wouters L, et al. "Graphical exploration of gene expression data: a comparative study of three multivariate methods", *Bioinformatics*. 2003 Dec;59(4):1131-9.
- [Yeang2001] Yeang, C., et al. "Molecular Classification of Multiple Tumor Types", In: *Proceedings of Ninth International Conference on Intelligent Systems in Molecular Biology*, Copenhagen, Denmark (July 21-25, 2001), S316-S322, 2001.
- [Yeo2001] Yeo G., Poggio, T. "Mutliclass Classification of SRBCT Tumors", Technical Report AI Memo 2001-018 CBCL Memo 206. MIT.