Published in final edited form as:

IEEE/ACM Trans Comput Biol Bioinform. 2012 ; 9(5): 1422–1431. doi:10.1109/TCBB.2012.63.

Gene selection using iterative feature elimination random forests for survival outcomes

Herbert Pang,

Biostatistics and Bioinformatics Department, Duke University School of Medicine, Durham, NC 27705. herbert.pang@duke.edu

Stephen L. George,

Biostatistics and Bioinformatics Department, Duke University School of Medicine, Durham, NC 27705. stephen.george@duke.edu

Ken Hui, and

School of Medicine, Yale University, New Haven, CT 06510. ken.hui@yale.edu

Tiejun Tong

Mathematics Department, Hong Kong Baptist University, Hong Kong SAR, China.
tongt@hkbu.edu.hk

Abstract

Although many feature selection methods for classification have been developed, there is a need to identify genes in high-dimensional data with censored survival outcomes. Traditional methods for gene selection in classification problems have several drawbacks. First, the majority of the gene selection approaches for classification are single-gene based. Second, many of the gene selection procedures are not embedded within the algorithm itself. The technique of random forests has been found to perform well in high dimensional data settings with survival outcomes. It also has an embedded feature to identify variables of importance. Therefore, it is an ideal candidate for gene selection in high dimensional data with survival outcomes. In this paper, we develop a novel method based on the random forests to identify a set of prognostic genes. We compare our method with several machine learning methods and various node split criteria using several real data sets. Our method performed well in both simulations and real data analysis. Additionally, we have shown the advantages of our approach over single-gene based approaches. Our method incorporates multivariate correlations in microarray data for survival outcomes. The described method allows us to best utilize the information available from microarray data with survival outcomes.

Index Terms

Cancer; Gene selection; Iterative feature elimination; Microarrays; Random Forest; Survival

Software

Three R packages `geneSelRSF`, `pgeneSelRSF` (parallelized version of `geneSelRSF`), and `geneSelCIF` are available for download at [62]: (<http://www.duke.edu/~hp44/geneSelRSF.html>).

1 Introduction

The identification of genes with high prognostic value and predictive power in clinical trials may help patients with cancer and other diseases. This may result from discovery of genes that serve as novel drug targets and/or those that can help guide better therapeutic decisions.

Many gene selection methods in the classification setting have been proposed in recent years. One of the earlier approaches was developed by Díaz-Uriarte and Alvarez de Andrés [1] for selecting genes with a random forests recursive feature elimination algorithm for class prediction. Duan et al. [2] proposed a recursive support vector machine gene selection algorithm. Tang, Zhang, and Huang [3] described a two-stage recursive feature elimination strategy for gene selection. Nijima and Okuno [4] extended the Laplacian linear discriminant analysis and developed a new algorithm for unsupervised feature selection. Mao and Tang [5] proposed a regularized recursive Mahalanobis separability measure for gene selection for classification problems. Mundra and Rajapakse [6] developed a filter based support vector machine recursive feature elimination based on mutual information for the same purpose. Luo et al. [7] proposed an improved support vector machine recursive cluster elimination that utilizes clusters to further improve gene selection performance for the classification setting.

Bøvelstad, Nygård and Borgan [8] have pointed out the importance of survival prediction. There are some implicit feature selection techniques such as boosting for survival outcomes [9]. Others developed Bayesian approach using model averaging [10]. However, not many methods of multivariate gene selection have been developed for censored survival data. Current approaches are mostly univariate approaches, and the number of genes is usually chosen with an arbitrary cutoff. The purpose of this paper is to introduce an iterative feature elimination algorithm for gene selection using random survival forests. The technique of random forests is among the best for correlating survival time with microarray data [11, 12]. We extend existing methods to survival phenotype and compare our method with other machine learning approaches that are known to do well in the survival setting. Our approach provides an alternative to univariate gene selection (such as Dunkler et al. [13]) for survival outcomes, with enhanced operating characteristics allowing the researchers to focus on a smaller set of genes for high dimensional data.

2 Methods

2.1 Random Forests

The random forest technique was first proposed for classification and regression settings [14]. The first version of survival forests, random forests for censored data, was implemented in Fortran [15]. The key difference between random forests classification/regression and its survival counterpart is that for random survival forests [16], the outcome of interest is a set of survival times with a corresponding censoring indicator. As a result, the split criteria for censored survival data are different from standard random forests for classification and regression, which uses the Gini criterion [14].

These differences are outlined in more detail in Section 2.2. However, the main properties of the original random forests are preserved in the random survival forests algorithm. For gene expression data, we treat each individual gene as a continuous intensity value. A random survival forest encompasses many binary trees, each formed by a deterministic algorithm. First, each tree is built using a bootstrap sample of the patients of interest. Second, a best binary split is chosen using a subset of genes in the data set. Unlike classification and regression trees (CART), no pruning is involved. Several split criteria are available in random survival forests; we apply log-rank, log-rank score, conserve, and random, for split

criteria as described in Section 2.3. Each observation is assigned to a leaf, the terminal node of a tree, according to the order and values of the predictor variables.

2.2 Random Survival Forests Algorithm

1. Draw bootstrap samples from the original data n_{tree} times, where n_{tree} is the number of trees. For each bootstrap sample, some observations are left out-of-bag (OOB).
2. A binary survival tree is grown for each bootstrap sample.
3. Let p be the number of genes in the data set; at each node of the tree, the number of genes selected at random for splitting will be $p^{1/2}$. This choice as suggested by Lunetta et al. does not have strong influence on the prediction error and variable importance [14, 15, 17].
4. Using one of the split criteria described in Section 2.3, a node is split using the single gene from the $p^{1/2}$ randomly chosen genes that maximizes the survival difference between the children nodes.
5. The splitting continues until each terminal node size reaches the minimum number of events with unique survival times. The default is three for right censored data [16].
6. The binary survival trees are then aggregated to obtain the ensemble cumulative hazard estimates as detailed below.

The binary trees are aggregated to form the forest through using ensemble cumulative hazard function (eCHF). The idea is essentially grouping the hazard estimates from the terminal nodes. The CHF estimates for a terminal node L is the Nelson-Aalen estimator (1):

$$\hat{\Lambda}_L(t) = \sum_{t_{i,L} \leq t} \frac{d_{i,L}}{R_{t_{i,L}}} \quad (1)$$

where $t_{i,L}$ = distinct survival time, $d_{i,L}$ = the number of events and $R_{t_{i,L}}$ = the number of individuals at risk at time $t_{i,L}$. For every binary survival tree with Q terminal nodes, there will be Q different CHF estimators. The CHF estimate for an individual i_{new} with gene predictor $gene_{new}$ can be found by identifying which terminal node includes the individual. That is, the CHF estimate is equal to $\hat{\Lambda}_L(t)$ if i_{new} ended in terminal node L . The eCHF is simply the sum of the CHFs across the bootstrap samples divided by number of trees. Let n be the total number of individuals, the expected number of ensemble events can be obtained by summing over time T_j for $j=1$ to n .

The OOB error rate is based on a prediction measure called Harrell's concordance index, denoted as C [18]. The OOB error rate, defined as $1 - C$, where $C \in [0,1]$, measures how well the random forests correctly ranks the survival of any pair of individuals. A value of $(1 - C)$ of 0.5 corresponds to a random guess and 0 means perfect concordance. The random survival forests implements both 'randomsplit' and 'permute' for variable importance. The latter is the same way Breiman used to perturb a variable in the original Random Forests algorithm [14]. The variable importance for each gene, used in Step 1) of Section 2.4, is calculated by permutation. The random survival forests algorithm permutes the values of the gene in the OOB cases, and the cases with permuted values are dropped down their in-bag survival tree. The cumulative hazard function is then calculated for each tree and aggregated across the trees. The randomly permuted values of the gene in the OOB individuals and the outcome of interest are independent of each other. The variable importance for a predictor x is equal to $PE_o - PE_n$, where PE_o is the prediction error of the original ensemble and PE_n is

the prediction error of the new ensemble. Variable importance is based on the concordance index described above and measures the extent of misclassification when the variable of interest becomes independent of outcome for the random survival forests.

2.3 Split Criteria

We will describe two split criteria that performed well in our study. Let i ($i = 1, \dots, n$) denotes a single individual; and x denotes one of the genes. The proposed split is of form $x \leq c$ and $x > c$, where c is the cutoff value.

Log-Rank—The log-rank (LR) split criterion (2) [19], which measures the node separation, is based on the log-rank test statistic defined as:

$$\text{LR}(X, c) = \frac{\sum_{i=1}^E d_{t_i, \text{child}_1} - E(D_i)}{\sqrt{\left[\sum_{i=1}^E \text{var}(D_i) \right]}} \quad (2)$$

where E is the number of distinct event times in the parent node, d_{t_i, child_j} is the number of events at time t_i in the child nodes $j=1, 2$, R_{t_i, child_j} is the number of individuals at risk at time t_i in the child nodes $j=1, 2$, and $R_{t_i} = \sum_{j=1}^2 R_{t_i, \text{child}_j}$, $d_{t_i} = \sum_{j=1}^2 d_{t_i, \text{child}_j}$, D_i is the random variable corresponding to the number of events in child node $j=1$ for the i -th distinct event time, $E(D_i) = R_{t_i, \text{child}_1} (d_{t_i} / R_{t_i})$ is the expectation, $\text{Var}(D_i) = d_{t_i} (R_{t_i} - d_{t_i}) / (R_{t_i} - 1) \times R_{t_i, \text{child}_1} / R_{t_i} \times (1 - R_{t_i, \text{child}_1} / R_{t_i})$ is the variance. The best split is defined as the one that maximizes the absolute value of the equation above.

Log-Rank Score (LRS)—Let $n_1 = \sum_{i=1}^n \mathbf{I}(X_i \leq c)$, where $\mathbf{I}(\cdot)$ is an indicator function. The log-rank score (LRS) split criterion (3), which measures the node separation, uses the following equation

$$\text{LRS}(X, c) = \frac{\sum_{X_i \leq c} a_i - n_1 \bar{a}}{\sqrt{n_1 \left(1 - \frac{n_1}{n}\right) s_a^2}} \quad (3)$$

where for individual i , $a_i = \mathbf{I}_i - \sum_{l=1}^{\gamma_i} \frac{\mathbf{I}_l}{n - \gamma_l + 1}$ as defined in Hothorn and Lausen [22], where $\mathbf{I}_i = 1$ if an event is observed for individual i and 0 otherwise; $\gamma_i = \sum_{l=1}^n \mathbf{I}(S_l \leq S_i)$ is the number of observed events or censored, occurring at survival time S_i or before; and $\bar{a}$ and s_a^2 are the sample mean and sample variance of a_i , respectively. The best split is defined as the one that maximizes the absolute value of the LRS equation above.

Conserve split criterion—Another type of splitting rule is the conservation of events [23]. Denote the Nelson-Aalen cumulative hazard estimator (4) for child j as:

$$\Lambda_j(t) = t_{i,j} \leq t \frac{d_{t_i, \text{child}_j}}{R_{t_i, \text{child}_j}} \quad (4)$$

where $t_{i,j}$ are the ordered event times for child j .

The conservation of events asserts that the total number of events is conserved in each child (5), i.e.

$$\sum_{i=1}^{n_j} \Lambda(T_{i,child_j}) = \sum_{i=1}^{n_j} 1_{i,child_j}(censoring) \quad (5)$$

where $1_{i,child_j}(censoring)$ is the censoring indicator.

Let $T_{(1),child_j}, T_{(2),child_j}, \dots, T_{(n),child_j}$ be the ordered time points for child j , and $1_{(i),child_j}(censoring)$ be the corresponding censoring indicator for $T_{(i),child_j}$ for $k=1, \dots, n_j$

$$Con(X, c) = \frac{1}{R_{t_1,child_1} + R_{t_1,child_2}} \sum_{j=1}^2 R_{t_1,child_j} \sum_{k=1}^{n_j-1} |M_{k,j}|, \quad (6)$$

where
$$M_{k,child_j} = \sum_{i=1}^k \hat{\Lambda}(T_{(i),child_j}) - \sum_{i=1}^k 1_{(i),child_j}(censoring)$$

The $Con(X, c)$ (6) measures whether the two groups are well separated and $(1 + Con(X, c))^{-1}$ is used in the program. It finds the best split by finding children closest to the conservation of events principle.

Random split criterion—The last splitting rule is "random" which implements a purely random uniform splitting [24]. For each node, a variable is randomly selected from a random set of p variables. For a chosen split variable X , a random split point is chosen among all possible split points on that variable.

2.4 Gene Selection with Random Survival Forests

The schema for the random survival forests gene selection algorithm is presented in Section A) of the supplementary materials. A description is as follows:

1. Gene expression and survival data are fed as input to the random survival forests algorithm. For each gene, $v_1, v_2, v_3, \dots, v_p$, the variable importance is calculated.
2. For iteration i , order the genes by variable importance in descending order and remove the bottom 20% (default). 20% is also the default chosen by [1] and real data analysis also suggests that this is a good choice, see Section 5.1. Call this removed set gs_i .
3. Calculate the OOB error rate ($1 - C$ as defined above) associated with the set gs_i using the expected number of ensemble events from random survival forests and the observed survival outcome.
4. Repeat Step 2 until gs_i contains only 2 genes. This is the limit for inclusion if we consider a multivariate gene selection.
5. Find the set of genes with the minimum number of genes such that the OOB error rate is within 1 standard error (s.e.), default, of the minimum OOB error rate, $mOOB$, of all iterations. i.e. $\text{argmin} \{ \#gs_i : mOOB - s.e. \leq OOB_i \leq mOOB + s.e. \}$. This one s.e. rule is commonly used in the classification tree literature [1, 23, 24]. The sampling error is determined by OOB_i and the total number of subjects.

2.5 Estimation of error rates

Given the iterative nature of the approach, the OOB error underestimates the true error rate and cannot be used to assess the overall error rate. This is analogous to the reasons that lead to selection bias [25, 26]. Therefore, we will use the 10 times 10-fold cross-validation, in which the gene selection procedure is applied to each training sample. The selected genes will then be used to predict the expected number of ensemble events in the testing set to assess prediction error rates. However, from the perspective of selecting the final optimal number of genes, incorporating OOB error rates to perform gene selection is not necessarily a bad procedure [1, 27].

3 Other machine learning methods

3.1 Random Conditional Inference Forests

Conditional inference forests [28] (R package party [29]) involves building many conditional inference trees under a binary recursive partitioning algorithm. The first step involves variable selection by testing a global null hypothesis of independence between any of the predictors and the survival outcome of interest. The algorithm terminates if the hypothesis cannot be rejected, otherwise, the predictor with the strongest association to the survival outcome will be selected. A binary split on the log-rank score (LRS) split criterion is implemented using the selected predictor. The steps above are then recursively repeated.

3.2 Survival Neural Network

The neural network for survival analysis [30] (R package survnnet [31]) utilizes the commonly used structure of a feed-forward network. A neural network consists of an input layer, hidden layers and an output layer. Each layer is composed of several neurons which receive output from neurons in the previous layer. The output is generated based on a specific transformation. Weights are applied to the input values between the layers. The goal of a neural network procedure is to seek values for the weights to fit the training data well. In the training procedure, the weights are iteratively updated to minimize the differences of the observed output and the generated output. This is usually done using a back propagation algorithm. We use parametric survival net, fitting a Weibull distribution to the data. The log hazard is modeled as a function of covariates, input, with a single linear output. Other settings include three hidden layers, Weibull shape parameter $p = \exp(0.1)$, and skip-layer units set to true.

3.3 Survival Cox Boosting

Given a set of outcome and predictors, the goal is to approximate the outcome with a function. Cox boosting (R package mboost [32]) proceeds as follows to optimize a loss function. First, initialize the function to a constant offset. Second, compute the residuals defined as the negative partial log-likelihood for Cox models. Third, fit the negative gradient vector by a base procedure, in this case, a component-wise univariate linear model as described in [33]. Fourth, update the function by taking a step of size 0 to 1 of the base learner. The second to fourth steps are repeated until a predetermined number of iterations are completed. The boosting algorithm is based on functional gradient descent. For more details see [34,35].

3.4 Survival Support Vector Machine

The goal of the support vector machine for survival (R package survpack [36]) is to find a non-linear hyperplane that separates the individuals who had events before time t and those without an event at that point. The non-linear relationship is obtained from a reproducing kernel Hilbert space with a Gaussian kernel of width 1. This procedure is repeated at every

event time. As in soft-margin support vector machines, the survival counterpart allows for cases when the hyperplanes do not exist by penalizing observations that lie on the other side of the margin. The minimization problem involves a sum that is based on the concordance index [37]. The penalty term is set to 1. The optimization problem is a positive-definite quadratic problem and is solved by using sequential minimal optimization algorithm.

3.5 Gene Selection with Random Conditional Inference Forests or Survival Neural Network

To compare with our approach described Section 2.4, Random Conditional Inference Forests is incorporated with Cox regression gene selection following a similar scheme for the iterative feature elimination algorithm. Similarly, this is done for Survival Neural Network. A description is as follows:

1. For each gene, univariate Cox regression is run.
2. For iteration i , order the genes by p-values in ascending order and remove the bottom 20%.
3. Calculate the concordance index associated with the set gs_i using the predicted survival time from Random Conditional Inference Forests or Survival Neural Network.
4. Repeat Step 2 until gs_i contains only 2 genes.
5. Find the set of genes with the minimum number of genes such that the error rate is within 2 s.e. Two is used because error rate tends to be the lowest with the largest number of genes chosen based on p-values from univariate Cox regression with either conditional inference forest or survival neural network.

4 Data Sets

4.1 Real Data Sets

The data sets used are given in Table 1. To improve the efficiency of the procedure, genes were filtered to contain only those that are mapped to known genes in cancer pathways [<http://www.broad.mit.edu>].

4.2 Simulation Data Sets

Our simulation design is to investigate the performance of our algorithm both under the null of no signal and under the alternative with different settings. This setup is similar to [1]. We simulated data under various settings with different numbers of independent dimensions, $d = 1, 2, 3$ and 10 , $b = 100$, and different number of genes per dimension, $g = 5, 20, 100$. The sample size n is 50 . The genes in each dimension are simulated from a multivariate normal distribution $MVN(0, \Sigma)$.

All genes have a variance of 1 from the multivariate normal distribution, and the correlation between genes within a dimension is 0.9, whereas the correlation between genes among dimensions is 0. In other words, the covariance matrix Σ is a block-diagonal matrix of size $d*b$ by $d*b$ with each of the b by b diagonal blocks Σ_p , and the rest of the matrix is zero. After the genes that belong to the dimensions are generated, we add another 1000 $U[-1, 1]$ variables to the original matrix of genes. The variance-covariance matrix is a block-diagonal matrix:

$$\Sigma = \begin{pmatrix} \Sigma_\rho & 0 & \dots & \dots & \vdots \\ 0 & \Sigma_\rho & 0 & \ddots & \vdots \\ \vdots & 0 & \Sigma_\rho & 0 & \vdots \\ \vdots & \ddots & 0 & \Sigma_\rho & 0 \\ \vdots & \dots & \dots & 0 & \Sigma_\rho \end{pmatrix}_{(d \times b) \times (d \times b)}$$

$$\Sigma_\rho = \begin{pmatrix} 1 & 0.9 & 0.9 & 0.9 & 0.9 \\ 0.9 & 1 & 0.9 & 0.9 & 0.9 \\ 0.9 & 0.9 & 1 & 0.9 & 0.9 \\ 0.9 & 0.9 & 0.9 & 1 & 0.9 \\ 0.9 & 0.9 & 0.9 & 0.9 & 1 \end{pmatrix}_{b \times b}$$

The survival time S , is defined as $S = \exp(-\mathbf{X}\beta + \epsilon)$ where $\mathbf{X}$ is the simulated gene expression matrix and β represents the coefficient parameter. For the null case, β equals 0 for all the predictors, $\epsilon \sim N(0,2)$, and the generated survival time is permuted for each run. Under the alternative case, β equals $1/2d$ for different numbers of informative genes, $i_g = 2, 4, 5, 6$, and 10, the ϵ term disappears, and $\beta = 0$ otherwise. The censoring time was generated from $N(\max(S), 20)$, which yields 20–40% censoring for each simulated data set. We have also investigated when $\beta = 1/3d$ and $1/d$; $\epsilon \sim N(0,1)$ and $N(0,3)$.

To calculate the error rate under the null and alternative cases, we first obtain the selected genes, gs , and calculate the 10-fold cross validation error rate and then repeat the 10-fold cross validation using a randomly selected set of genes of $\#gs$.

5 Results and Discussions

5.1 Application to Data Sets

MCL is a data set that consists of microarray and survival data for patients with mantle cell lymphoma [38]. DCL is a data set that contains microarrays and survival information for patients with diffuse large B-cell lymphoma after chemotherapy [39]. NKI is a data set that includes microarray and survival data for patients with primary breast cancers [40].

Table 2 shows a comparison of different machine learning algorithms for survival prediction and random survival forests with gene selection (top half) or without gene selection, i.e. using all genes (bottom half). The average 10 times 10-fold cross validation error rates from the algorithms with gene selection is comparable to those using all genes. This finding is consistent with Díaz-Uriarte and Alvarez de Andrés, who investigated gene selection in random forests classification [1]. The numbers of genes selected by the different algorithms are displayed in Table 3. Table 3 also shows the number of genes selected in the different training sets, pairwise correlation among the selected genes (original), and the proportion of genes selected in the original set that appears in the training set. Unlike the cforest+selection and neural network+selection, our iterative feature elimination algorithm with random survival forests has lower correlation among genes as it is a multivariate-based gene selection algorithm. In terms of proportions of genes in training that are in the original data, cforest+selection and our iterative feature elimination algorithm performed similarly. Even though both cforest+selection and neural network+selection uses univariate Cox model as the method for ranking genes, the iterative feature elimination algorithm for the former

outperforms the latter in all three data sets. Integrated Brier Score, which measures how well the survival function is estimated, was also used, see Section B) of the supplementary materials.

Table 2 shows the comparison of different split methods. Overall, none of the split criteria from random survival forests is superior to others. Given the variation in performance, unless there is a special preference, we recommend log-rank for the iterative elimination algorithm as it is the default chosen for random survival forests [16]. Table 4 shows how the gene selection algorithm performs under an alternative choice of 0 and 0.5 s.e. In terms of average CV error, the default 1 s.e. and the alternative choices of standard errors results are very similar. However, the number of genes selected are 168, 58, and 77 for MCL, DCL and NKI, respectively for 0 s.e. The high number of genes is much less desirable than what we found in the case of 1 s.e., with 3, 15, and 14 genes, respectively. The one standard error rule is consistent with what has been used in the machine learning literature [23, 24]. 0.5 s.e. which produces higher proportion of genes selected in the original for DCL and NKI may be considered if higher number of genes is desired. Figure of Section C) in supplementary materials is a plot of the average 10 times 10-fold cross validation error rates against the number of trees used to perform the gene selection.

For all three data sets, the average cross validation error rates are consistent for 1,000 trees or more. Therefore, 1,000 trees are used in all other analyses. For the MCL data set, the error rate converges with around 500 trees. Figure of Section D) in supplementary materials investigates the tradeoff between increasing correlation or strength of the trees among a forest. Overall, it shows for the three data sets that the default number of genes used for splitting performs quite well. Table 5 illustrates the impact of the percentage of features dropped on the average cross validation error rates. If we increase the percentage of features dropped from the 20% to 50%, the average cross validation error rates are slightly worse with 5, 21, and 11 genes chosen for MCL, DCL and NKI, respectively. And for 80% dropped, it essentially removes the flexibility of the iterative feature elimination algorithm because the number of genes is fixed for all runs at 10, 11, and 22 for MCL, DCL, and NKI, respectively. Therefore, we recommend dropping 20% of the features at each iteration with 1000 trees.

The selected genes for MCL were PRIM1, PCNA, and ASPM. PCNA has been shown to be associated with proliferation and survival of mantle cell lymphoma [41, 42]. The selected genes for DCL were PHKB, FLJ20703, BMP6 (2x), PTGES, TKT, PLAU, SPARC, CCL2, CXCR6, ITGAL, DAXX, CTSE, NFKBIL1, and SH3BGRL. BMP6, which was selected twice, has been identified as a potential risk predictor for diffuse large B-cell lymphoma (DLBCL) and is hypermethylated in aggressive forms of lymphoma [43]. SPARC is one of the genes in a CD5 DLBCL signature [44]. CCL2, also known as MCP-1, is produced in primary central nervous system lymphomas which consist mostly of DLBCL [45], and it is involved in the migration and localization of follicular lymphoma cells [46]. CXCR6 is one of the four chemokine receptors identified as a good classifier between MALT lymphoma and extranodal DLBCL [47]. ITGAL, also known as CD11a, is detected on some B-cell clones and is hypothesized to be related to tissue localization of B-Chronic Lymphocytic Leukemia [48]. Two research groups have discovered that DAXX is closely tied with the disease mechanism of B-cell lymphomas [49, 50]. A common NFKBIL1 variant is found to be associated with non-Hodgkin lymphoma risk [51]. The selected genes for NKI are DNMT3B, FAF1, RFC4, ADFP, CIZ1, PIR, ABCF1, CANX, AP3S2, SERPINA5, CTSL2, TMEPAI, CCNB1, and HMOX1. CIZ1 is an estrogen receptor coactivator, and plays a regulatory role in the cell cycle progression of breast cancer [52, 53]. Researchers using derived PIR, PARP-inhibitor-resistant, clones were able to show that resistance to PARP inhibition can be acquired by deletion of a mutation in BRCA2 which is a well known breast

cancer gene [54]. Other researchers have found SERPINA5 to be a potential positive prognostic factor for breast cancer [55]. CCNB1 has been identified as a prognostic marker in lymph node negative breast cancer [56]. Further discussions of these genes can be found in Section E) of supplementary materials.

We also compared our proposed algorithm with that of univariate gene selection (single-gene based) using Cox regression and correcting for multiple comparison using FDR [57, 58]. Under this method, the numbers of genes found to be significant at 0.05 q-value levels were 497, 7 and 897 for MCL, DCL and NKI, respectively. For the NKI data set, there were 897 found by Cox model with 5% FDR, and only two of the 14 found by our method did not overlap with Cox model. For the DCL data set, there were 7 found by Cox model with 5% FDR, only two overlapping with our method. For MCL, there were 497 found by the Cox model with 5% FDR, and only one of the three found by our method did not overlap with the Cox model. We performed DAVID analysis to look at enriched pathways, and the DCL and NKI data sets did not result in any significantly enriched pathways after multiple testing adjustments [59]. For the MCL data set, 15 KEGG pathways were found to be significant. When we use ($SE = 0$), it identified more genes, as it chooses the iteration that has the lowest error rate. For NKI, 24 of the 77 were not found in the 897 from the univariate Cox regression model with 5% FDR. For DCL, 2 of the 58 were not found in the 7 from the univariate Cox regression model with 5% FDR. And for MCL, 107 of the 168 were not found in the 497 from the univariate Cox regression model with 5% FDR. With the exception of DCL, the set of genes found by using our approach is much more manageable and picks out a realistic set of genes for further testing and validation. In addition, our approach uses random forests which implicitly takes into account the way the genes interact.

Unlike univariate gene selection, our approach does not require the user to set a threshold cutoff for p-values or q-values. Our algorithm contains tuning parameters such as the s.e. that can be modified to vary the number of genes selected. We recommend using 1 s.e. as it is commonly used in the classification tree setting. Additionally, a reasonable number of genes were obtained from the three real data sets analyzed. These genes can be tested for further validation.

5.2 Simulation Results

To assess the cross validation error rate under the null of no signal and the alternative, we performed simulations as described above. For every simulated data set, we first performed the gene selection procedure and then obtained the error rates from the 10-fold cross-validation using the selected genes. For both type I error and power, we calculated the average of the 10-fold cross-validation error over 100 simulations.

From Table 6, under the case of no signal, we see that the average cross-validation error rate ranges from 0.43 to 0.45, slightly better than random. The median and average number of genes chosen is approximately two to three for various settings, including genes per dimension of 5, 20, 50 and 100 and dimensions of 1, 2, 3 and 10. For the case with informative genes, Table 7, we see that the median and average number of genes selected across the simulations is very close to the true value under different settings. The error rates are close to 0.05 for the one-dimension cases with various numbers of genes per dimension. As the number of dimensions increases to two and three, the error rates hover around 0.11 to 0.15. The above simulation demonstrates that our method is performing well under the null and alternative cases. Additional results can be found in Section F) and G) of supplementary materials.

5.3 Computational Time

For a comparison of the computational time of the single processor and the parallelized versions, please see Table 8. Parallelization takes place after the importance measure was calculated for the data set as a whole, therefore, only the error rate calculations were parallelized. Performance can be further improved if the parallelization can take place within the C code called within R.

6 Conclusions

High-throughput genomics data can help predict the prognosis of patients in complex diseases like cancer. We developed a method to identify a set of prognostic genes that predicts the survival of patients, compared several machine learning methods and various node split criteria using several existing data sets, and evaluated our method via simulation. Our method performed well under the null and alternative cases. We applied our method to data sets on breast cancer, B-cell lymphoma, and mantle cell lymphoma. Using these data sets, we demonstrated the disadvantages of using a univariate gene selection method. Our approach has the advantage of being able to identify a small set of genes while preserving the predictive accuracy for survival. Univariate methods which do not take correlation into account can often pick out too many genes, making further validation difficult. However, they can be coupled with iterative feature elimination algorithms like cforest+selection developed in this paper. We also investigated the performance of our gene selection algorithm under various settings, including the percentage of features dropped, the standard error for choosing mOOB, and the number of trees used. To the best of our knowledge, our gene selection method is the first proposed to incorporate multivariate correlations in microarray data for survival outcome using machine learning methods. It will be interesting to compare the proposed method with other approaches based on gene pairs [60]. We emphasize that the goal of this procedure is not to identify signature, but to pick out genes for further validation. The algorithm proposed in this paper is for gene selection for a particular training set in high dimensional data, such as microarrays. In the future, incorporating pathways into gene selection should help enhance the predictive ability and interpretability of the findings [61].

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

Funding: National Institutes of Health (grant P01CA142538) and start-up funds from Duke University Medical Center.

References

1. Díaz-Uriarte R, Alvarez de Andrés S. Gene selection and classification of microarray data using random forest. *BMC Bioinformatics*. 2006; vol. 7:3.
2. Duan KB, Rajapakse JC, Wang H, Azuaje F. Multiple SVM-RFE for gene selection in cancer classification with expression data. *IEEE Trans Nanobioscience*. 2005; vol. 4(no. 3):228–234. [PubMed: 16220686]
3. Tang Y, Zhang YQ, Huang Z. Development of two-stage SVM-RFE gene selection strategy for microarray expression data analysis. *IEEE/ACM Trans Comput Biol Bioinform*. 2007; vol. 4(no. 3): 365–381. [PubMed: 17666757]
4. Nijima S, Okuno Y. Laplacian linear discriminant analysis approach to unsupervised feature selection. *IEEE/ACM Trans Comput Biol Bioinform*. 2009; vol. 6(no. 4):605–614. [PubMed: 19875859]

5. Mao KZ, Tang W. Recursive Mahalanobis Separability Measure for Gene Subset Selection. *IEEE/ACM Trans Comput Biol Bioinform.* Epub. May 10, 2010.
6. Mundra PA, Rajapakse JC. SVM-RFE with MRMR filter for gene selection. *IEEE Trans Nanobioscience.* 2010; vol. 9(no. 1):31–37. [PubMed: 19884101]
7. Luo LK, Huang DF, Ye LJ, Zhou QF, Shao GF, Peng H. Improving the Computational Efficiency of Recursive Cluster Elimination for Gene Selection. *IEEE/ACM Trans Comput Biol Bioinform.* Epub. May 14, 2010.
8. Bøvelstad HM, Nygård S, Borgan O. Survival prediction from clinico-genomic models--a comparative study. *BMC Bioinformatics.* 2009; vol. 10:413.
9. Binder H, Schumacher M. Allowing for mandatory covariates in boosting estimation of sparse high-dimensional survival models. *BMC Bioinformatics.* 2008; Vol. 10:9–14.
10. Lee K, Mallick B. Bayesian methods for variable selection in survival models with application to DNA microarray data. *Sankhya.* 2004; Vol. 66:756–778.
11. Schumacher M, Binder H, Gerds T. Assessment of survival prediction models based on microarray data. *Bioinformatics.* 2007; vol. 23:1768–1774. [PubMed: 17485430]
12. van Wieringen W, Kun D, Hampel R, Boulesteix A. Survival prediction using gene expression data: a review and comparison. *Comput Stat Data Anal.* 2009; vol. 53:1590–1603.
13. Dunkler D, Schemper M, Heinze G. Gene selection in micro-array survival studies under possibly non-proportional hazards. *Bioinformatics.* 2010; vol. 26:784–790. [PubMed: 20118118]
14. Breiman L. Random forests. *Mach. Learn.* 2001; vol. 45:5–32.
15. Breiman L. How to use survival forests (SFPDV1). [http://oz.berkeley.edu/users/breiman/SF_Manual.pdf].
16. Ishwaran H, Kogalur U, Blackstone E, Lauer M. Random survival forests. *Ann Appl Stat.* 2008; vol. 2:841–860.
17. Lunetta KL, Hayward LB, Segal J, Van Eerdewegh P. Screening large-scale association study data: exploiting interactions using random forests. *BMC Genet.* 2004; vol. 5:32. [PubMed: 15588316]
18. Harrell FE, Califf RM, Pryor DB, Lee KL, Rosati RA. Evaluating the yield of medical tests. *Journal of the American Medical Association.* 1982; vol. 247:2543–2546. [PubMed: 7069920]
19. Segal M. Regression trees for censored data. *Biometrics.* 1988; vol. 44:35–47.
20. Hothorn T, Lausen B. On the exact distribution of maximally selected rank statistics. *Comput Stat Data Anal.* 2003; vol. 43:121–137.
21. Naftel D, Blackstone E, Turner M. Conservation of events. 1985 unpublished.
22. Lin Y, Jeon Y. Random forests and adaptive nearest neighbors. *J. Amer. Stat. Assoc.* 2006; 101:578–590.
23. Breiman, L.; Friedman, J.; Olshen, R.; Stone, C. Classification and regression trees. New York: Chapman & Hall; 1984.
24. Ripley, B. Pattern recognition and neural networks. Cambridge University Press; 1996.
25. Ambroise C, McLachlan GJ. Selection bias in gene extraction on the basis of microarray gene-expression data. *PNAS.* 2002; vol. 99:6562–6566. [PubMed: 11983868]
26. Simon R, Radmacher M, Dobbin K, McShane L. Pitfalls in the use of DNA microarray data for diagnostic and prognostic classification. *JNCI.* 2003; vol. 95:14–18. [PubMed: 12509396]
27. Braga-Neto U, Hashimoto R, Dougherty E, Nguyen D, Carroll R. Is cross-validation better than resubstitution for ranking genes? *Bioinformatics.* 2004; vol. 20:253–258. [PubMed: 14734317]
28. Hothorn T, Hornik K, Zeileis A. Unbiased recursive partitioning: a conditional inference framework. *J Comput Graph Stat.* 2006; vol. 15:651–674.
29. Strobl C, Boulesteix A-L, Zeileis A, Hothorn T. Bias in Random Forest Variable Importance Measures: Illustrations, Sources and a Solution. *BMC Bioinformatics.* 2007; vol. 8
30. Ripley R, Harris A, Tarassenko L. Non-linear survival analysis using neural networks. *Stat Med.* 2004; vol. 23:825–842. [PubMed: 14981677]
31. Ripley R. survnnet: Feed-forward neural networks for survival analysis. R package version 1.1–2.
32. Hothorn T, Buhlmann P, Kneib T, Schmid M, Hofner B. mboost: Model-Based Boosting. R package version 1.0–5. 2008

33. Buhlmann P, Hothorn T. Boosting algorithms: regularization, prediction and model fitting. *Stat Sci.* 2007; 22:477–505.
34. Friedman J. Greedy function approximation: a gradient boosting machine. *Annals of Statist.* 2001; 29:1189–1232.
35. Bühlman P, Yu B. Boosting with the L2 loss: regression and classification. *Journal of the American Statistical Association.* 2003; vol. 98:324–339.
36. Evers L. survpack: Methods for Fitting High-Dimensional Survival Models. R package version 0.1–4. 2008
37. Evers L, Messow C. Sparse kernel methods for highdimensional survival data. *Bioinformatics.* 2008; 15:1632–1638. [PubMed: 18515276]
38. Rosenwald A, Wright G, Wiestner A, Chan WC, Connors JM, Campo E, Gascoyne RD, Grogan TM, Muller-Hermelink HK, Smeland EB, Chiorazzi M, Giltane JM, Hurt EM, Zhao H, Averett L, Henrickson S, Yang L, Powell J, Wilson WH, Jaffe ES, Simon R, Klausner RD, Montserrat E, Bosch F, Greiner TC, Weisenburger DD, Sanger WG, Dave BJ, Lynch JC, Vose J, Armitage JO, Fisher RI, Miller TP, LeBlanc M, Ott G, Kvaloy S, Holte H, Delabie J, Staudt LM. The proliferation gene expression signature is a quantitative integrator of oncogenic events that predicts survival in mantle cell lymphoma. *Cancer Cell.* 2003; vol. 3:185–197. [PubMed: 12620412]
39. Rosenwald A, Wright G, Chan WC, Connors JM, Campo E, Fisher RI, Gascoyne RD, Muller-Hermelink HK, Smeland EB, Giltane JM, Hurt EM, Zhao H, Averett L, Yang L, Wilson WH, Jaffe ES, Simon R, Klausner RD, Powell J, Duffey PL, Longo DL, Greiner TC, Weisenburger DD, Sanger WG, Dave BJ, Lynch JC, Vose J, Armitage JO, Montserrat E, López-Guillermo A, Grogan TM, Miller TP, LeBlanc M, Ott G, Kvaloy S, Delabie J, Holte H, Krajci P, Stokke T, Staudt LM. Lymphoma/Leukemia Molecular Profiling Project: The use of molecular profiling to predict survival after chemotherapy for diffuse large-B-cell lymphoma. *N Engl J Med.* 2002; vol. 346:1937–1947. [PubMed: 12075054]
40. van de Vijver MJ, He YD, van't Veer LJ, Dai H, Hart AA, Voskuil DW, Schreiber GJ, Peterse JL, Roberts C, Marton MJ, Parrish M, Atsma D, Witteveen A, Glas A, Delahaye L, van der Velde T, Bartelink H, Rodenhuis S, Rutgers ET, Friend SH, Bernards R. A gene-expression signature as a predictor of survival in breast cancer. *N Engl J Med.* 2002; vol. 347:1999–2009. [PubMed: 12490681]
41. Castillo R, Mascarenhas J, Telford W, Chadburn A, Friedman S, Schattner E. Proliferative response of mantle cell lymphoma cells stimulated by CD40 ligation and IL-4. *Leukemia.* 2000; vol. 14:292–298. [PubMed: 10673747]
42. Desai S, Maurin M, Smith M, Bolick S, Dessureault S, Tao J, Sotomayor E, Wright K. PRDM1 is required for mantle cell lymphoma response to bortezomib. *Mol Cancer Res.* 2010; vol. 8:907–918. [PubMed: 20530581]
43. Daibata M, Nemoto Y, Bandobashi K, Kotani N, Kuroda M, Tsuchiya M, Okuda H, Takakuwa T, Imai S, Shuin T, Taguchi H. Promoter hypermethylation of the bone morphogenetic protein-6 gene in malignant lymphoma. *Clin Cancer Res.* 2007; vol. 13:3528–3535. [PubMed: 17575215]
44. Suguro M, Tagawa H, Kagami Y, Okamoto M, Ohshima K, Shiku H, Morishima Y, Nakamura S, Seto M. Expression profiling analysis of the CD5+ diffuse large B-cell lymphoma subgroup: development of a CD5 signature. *Cancer Sci.* 2006; vol. 97:868–874. [PubMed: 16856881]
45. Kitai R, Ishisaka K, Sato K, Sakuma T, Yamauchi T, Imamura Y, Matsumoto H, Kubota T. Primary central nervous system lymphoma secretes monocyte chemoattractant protein 1. *Med Mol Morphol.* 2007; vol. 40:18–22. [PubMed: 17384985]
46. Husson H, Carideo E, Cardoso A, Lugli S, Neuberg D, Munoz O, de Leval L, Schultze J, Freedman A. MCP-1 modulates chemotaxis by follicular lymphoma cells. *Br J Haematol.* 2001; vol. 115:554–562. [PubMed: 11736935]
47. Deutsch A, Aigelsreiter A, Steinbauer E, Frühwirth M, Kerl H, Beham-Schmid C, Schaidner H, Neumeister P. Distinct signatures of B-cell homeostatic and activation-dependent chemokine receptors in the development and progression of extragastric MALT lymphomas. *J Pathol.* 2008; vol. 215:431–444. [PubMed: 18561120]

48. Kimby E, Rincon J, Patarroyo M, Mellsted H. Expression of adhesion molecules CD11/CD18 (Leu-CAMs, β 2-integrins), CD54 (ICAM-1) and CD58 (LFA-3) in B-chronic lymphocytic leukemia. *Leuk Lymph*. 1994; vol. 13:297.
49. Perlman R, Schiemann W, Brooks M, Lodish H, Weinberg R. TGF-beta-induced apoptosis is mediated by the adapter protein Daxx that facilitates JNK activation. *Nat Cell Biol*. 2001; vol. 3:708–714. [PubMed: 11483955]
50. Brieger A, Boehrer S, Schaaf S, Nowak D, Ruthardt M, Kim S, Atadja P, Hoelzer D, Mitrou P, Weidmann E, Chow K. In bcr-abl-positive myeloid cells resistant to conventional chemotherapeutic agents, expression of Par-4 increases sensitivity to imatinib (STI571) and histone deacetylase-inhibitors. *Biochem Pharmacol*. 2004; vol. 68:85–93. [PubMed: 15183120]
51. Wang SS, Purdue MP, Cerhan JR, Zheng T, Menashe I, Armstrong BK, Lan Q, Hartge P, Krickler A, Zhang Y, Morton LM, Vajdic CM, Holford TR, Severson RK, Grulich A, Leaderer BP, Davis S, Cozen W, Yeager M, Chanock SJ, Chatterjee N, Rothman N. Common gene variants in the tumor necrosis factor (TNF) and TNF receptor superfamilies and NF-kB transcription factors and non-Hodgkin lymphoma risk. *PLoS One*. 2009; 4(4):e5360. Epub 2009 Apr 24. [PubMed: 19390683]
52. den Hollander P, Rayala S, Coverley D, Kumar R. Ciz1, a Novel DNA-binding coactivator of the estrogen receptor alpha, confers hypersensitivity to estrogen action. *Cancer Res*. 2006; vol. 66:11021–11029. [PubMed: 17108141]
53. den Hollander P, Kumar R. Dynein light chain 1 contributes to cell cycle progression by increasing cyclin-dependent kinase 2 activity in estrogen-stimulated cells. *Cancer Res*. 2006; vol. 66:5941–5949. [PubMed: 16740735]
54. Edwards S, Brough R, Lord C, Natrajan R, Vatcheva R, Levine D, Boyd J, Reis-Filho J, Ashworth A. Resistance to therapy caused by intragenic deletion in BRCA2. *Nature*. 2008; vol. 451:1111–1115. [PubMed: 18264088]
55. Castelló R, Landete J, España F, Vázquez C, Fuster C, Almenar S, Ramón L, Radtke K, Estellés A. Expression of plasminogen activator inhibitors type 1 and type 3 and urokinase plasminogen activator protein and mRNA in breast cancer. *Thromb Res*. 2007:753–762.
56. Niméus-Malmström E, Koliadi A, Ahlin C, Holmqvist M, Holmberg L, Amini R, Jirström K, Wärnberg F, Blomqvist C, Fernö M, Fjällskog M. Cyclin B1 is a prognostic proliferation marker with a high reproducibility in a population-based lymph node negative breast cancer cohort. *Int J Cancer*. 2010; vol. 127:961–967.
57. Cox D. Regression Models and Life-Tables. *Journal of the Royal Statistical Society Series B*. 1972; vol. 34:187–220.
58. Storey J, Tibshirani R. Statistical significance for genome-wide studies. *PNAS*. 2003; vol. 100:9440–9445. [PubMed: 12883005]
59. Huang D, Sherman B, Lempicki R. Systematic and integrative analysis of large gene lists using DAVID Bioinformatics Resources. *Nature Protoc*. 2009; 4:44–57. [PubMed: 19131956]
60. Motakis E, Ivshina A, Kuznetsov V. Data-driven approach to predict survival of cancer patients. *IEEE Engineering in Medicine and Biology*. 2009; 28:58–66.
61. Pang H, Datta D, Zhao H. Pathway analysis using random forests with bivariate node-split for survival outcomes. *Bioinformatics*. 2010; vol. 26:250–258. [PubMed: 19933158]
62. R Development Core Team. Vienna, Austria: 2010. R: A language and environment for statistical computing. R Foundation for Statistical Computing. ISBN 3-900051-07-0, URL <http://www.R-project.org>

Biographies

Herbert Pang received a B.A. degree in Mathematics and Computer Science from University of Oxford, England, and a Ph.D. degree in Biostatistics from Yale University. He is currently an Assistant Professor in the Department of Biostatistics and Bioinformatics at Duke University. His main research interests include bioinformatics, genomics, and clinical trials.

Stephen L. George received a B.A. degree in Mathematics from Texas Tech University, a M.E.S. degree in Statistics from NC State University, and a Ph.D. degree in Statistics from Southern Methodist University. He is currently a Professor in the Department of Biostatistics and Bioinformatics at Duke University. He is a Fellow of American Statistical Association, a Fellow and former President of the Society of Clinical Trials, a former Executive Committee member of the International Biometric Society, and a former member of the FDA Oncology Drugs Advisory Committee. His research interests are in statistical methodology for the design and analysis of clinical trials.

Ken Hui received his B.S. and M.S. in Applied Mathematics and Biology from Yale University. He is currently an M.D./Ph.D. student at Yale University School of Medicine.

Tiejun Tong received his B.S. in Electronic Engineering and M.S. in Statistics from the University of Science and Technology of China, and a Ph.D. degree in Statistics from the University of California at Santa Barbara. He is currently an Assistant Professor in the Department of Mathematics at the Hong Kong Baptist University. His main research interests include high dimensional data analysis and shrinkage estimation.

\$watermark-text

\$watermark-text

\$watermark-text

Data Sets

TABLE 1

Data set	Reference	Sample Size	Genes in Cancer Pathways	All Genes	Response
MCL	Rosenwald et al. (2003)	92	5969	8810	Overall Survival
DCL	Rosenwald et al. (2002)	240	1320	7399	Overall Survival
NKI	van de Vijver et al. (2002)	295	2740	9746	Overall Survival

TABLE 2**Gene Selection Results**

Methods	MCL	DCL	NKI
<i>Selected genes</i>			
Random survival forests IFE log-rank split (w. selection)	0.34	0.44	0.29
Random survival forests IFE log-rank score split (w. selection)	0.39	0.39	0.29
Random survival forests IFE random split (w. selection)	0.34	0.45	0.29
Random survival forests IFE conserve split (w. selection)	0.35	0.42	0.29
Cforest+selection	0.29	0.43	0.29
Neural Network+selection	0.36	0.44	0.32
<i>All genes</i>			
Random survival forests log-rank split	0.32	0.42	0.26
Random survival forests log-rank score split	0.34	0.37	0.28
Random survival forests random split	0.39	0.41	0.28
Random survival forests conserve split	0.32	0.39	0.25
Conditional inference forests	0.34	0.4	0.23
Neural network	0.4	0.42	0.31
Support vector machine	0.42	0.42	0.26
Cox boosting	0.3	0.38	NA

Comparison of different methods with selected genes and all genes – average over 10 times (10-fold) cross validation error rates. 'NA' is reported for Cox Boosting for NKI data because an R error message occurred while running this algorithm for this data set.

TABLE 3**Summary Statistics of Gene Selection**

RSF IFE log-rank split	MCL	DCL	NKI
Error Rate	0.34	0.44	0.29
Genes selected (original)	3	15	14
Genes selected (training)	6 (4, 9)	10 (8, 12)	9 (9, 14)
Pairwise Correlation average	0.47	0.08	0.13
Proportions of genes selected in original set that appear in training set	33% (0, 33)	20% (13, 27)	21% (14, 29)
Cforest+selection	MCL	DCL	NKI
Error Rate	0.29	0.43	0.29
Genes selected (original)	5	29	12
Genes selected (training)	7 (4, 18)	9 (6, 12)	10 (8, 16)
Pairwise Correlation average	0.68	0.36	0.65
Proportions of genes selected in original set that appear in training set	20% (0, 60)	10% (7, 25)	25% (17, 42)
Neural Network+selection	MCL	DCL	NKI
Error Rate	0.36	0.44	0.32
Genes selected (original)	2	7	8
Genes selected (training)	3 (2, 5)	6 (4, 7)	6 (5, 8)
Pairwise Correlation average	0.77	0.10	0.12
Proportions of genes selected in original set that appear in training set	0% (0, 50)	14% (0, 28)	25% (13, 38)

Summary statistics of gene selection over the 10 times 10-fold cross validation run. Gene selected (training) – median (1st quartile, 3rd quartile).

Proportions of genes - median (1st quartile, 3rd quartile)

TABLE 4**Impact of Standard Error**

RSF IFE log-rank split Method (SE = 0)	MCL	DCL	NKI
Error Rate	0.33	0.41	0.27
Genes selected (original)	168	58	77
Genes selected (training)	55 (22, 86)	37 (24, 46)	32 (21, 50)
Proportions of genes selected in original set that appear in training set	8% (4, 11)	24% (19, 29)	17% (13, 21)
RSF IFE log-rank split Method (SE = 0.5)	MCL	DCL	NKI
Error Rate	0.34	0.43	0.28
Genes selected (original)	11	19	9
Genes selected (training)	14 (9, 22)	15 (12, 19)	14 (11, 18)
Proportions of genes selected in original set that appear in training set	0% (0, 9)	31% (26, 37)	44% (33, 56)

Impact on the average CV error rate under different standard error values for determining the mOOB

TABLE 5**Impact of Percent of Features Dropped**

RSF IFE log-rank split Method	MCL	DCL	NKI
(% Dropped=20) Error Rate	0.34	0.44	0.29
(% Dropped=50) Error Rate	0.35	0.43	0.31
(% Dropped=80) Error Rate	0.33	0.43	0.28
RSF IFE log-rank split Method (% Dropped=50)	MCL	DCL	NKI
Error Rate	0.35	0.43	0.31
Genes selected (original)	5	21	11
Genes selected (training)	5 (5, 11)	11 (11, 11)	11 (11, 21)
Proportions of genes selected in original set that appear in training set	20% (0, 20)	19% (19, 25)	27% (18, 36)
RSF IFE log-rank split Method (% Dropped=80)	MCL	DCL	NKI
Error Rate	0.33	0.43	0.28
Genes selected (original)	10	11	22
Genes selected (training)	10 (10, 10)	11 (11, 11)	22 (22, 22)
Proportions of genes selected in original set that appear in training set	0% (0, 10)	36% (27, 45)	32% (27, 36)

Impact on the average CV error rate with different percentage of features dropped

TABLE 6

Under The Null - No Signal

Dimensions	Genes per dimension	Average # of genes	Median # of genes	Average CV error rate
1	5	2.01	2	0.45
1	20	2.03	2	0.44
2	5	2.08	2	0.44
2	20	2.25	2	0.43
3	5	2.12	2	0.44
3	20	2.18	2	0.43
10	50	3.08	2	0.44
10	100	3.15	2	0.44

Simulation results: Average CV error rate under the null of no signal

\$watermark-text

\$watermark-text

\$watermark-text

TABLE 7

Under The Alternative - With Informative Genes

Dimensions	Genes per dimension	Informative genes	Average # of genes	Median # of genes	Average CV error rate
1	5	2	2.76	2	0.05
1	20	5	4.99	4	0.047
1	100	10	9.31	8	0.048
2	5	4	4.34	3	0.127
2	20	4	6.12	4	0.118
2	100	4	5.64	3	0.117
3	5	6	6.56	4	0.152
3	20	6	7.37	5	0.143
3	100	6	7.55	5	0.131

Simulation results: Average CV error rate under the alternative with informative genes

\$watermark-text

\$watermark-text

\$watermark-text

TABLE 8

Computational Time

	Single processor version	Cluster parallel version	
	(time in minutes)		
		7 nodes	14 nodes
MCL	20	> 30	n/a
DCL	59	47	36
NKI	116	86	76

Computational time for the three real data sets in minutes