

Prediction Error Estimation: A Comparison of Resampling Methods

Annette M. Molinaro *

Biostatistics Branch

Division of Cancer Epidemiology and Genetics, NCI, NIH

Richard Simon

Biometric Research Branch

Division of Cancer Treatment and Diagnostics NCI, NIH

Ruth M. Pfeiffer

Biostatistics Branch

Division of Cancer Epidemiology and Genetics, NCI, NIH

June 6, 2005

Abstract

In genomic studies, thousands of features are collected on relatively few samples. One of the goals of these studies is to build classifiers to predict the outcome of future observations. There are three inherent steps to this process: feature selection, model selection, and prediction assessment. With a focus on prediction assessment, we compare several methods for estimating the 'true' prediction error of a prediction model in the presence of feature selection. For small studies where features are selected from thousands of candidates, the resubstitution and simple split-sample estimates are seriously biased. In these small samples, leave-one-out (LOOCV), 10-fold cross-validation (CV), and the .632+ bootstrap have the smallest bias for diagonal discriminant analysis, nearest neighbor, and classification trees. LOOCV and 10-fold CV have the smallest bias for linear discriminant analysis. Additionally, LOOCV, 5- and 10-fold CV, and the .632+ bootstrap have the lowest mean square error. The .632+ bootstrap is quite biased in small sample sizes with strong signal to noise ratios. The differences in performance among resampling methods are reduced as the number of specimens available increases.

Supplementary Information: R code for simulations and analyses is available from the authors. Tables and figures for all analyses are available at <http://linus.nci.nih.gov/brb/TechReport.htm>

Keywords: Cross-validation, resampling, bootstrap, prediction, risk-estimation, feature selection, proteomics data, microarrays, small sample, class prediction.

*annette.molinaro@yale.edu

Contents

1	Introduction	3
2	Methods	4
2.1	Resampling Methods	5
2.1.1	Split Sample	5
2.1.2	v -fold Cross-Validation	5
2.1.3	Leave-One-Out Cross-Validation (LOOCV)	6
2.1.4	Monte Carlo Cross-Validation (MCCV)	6
2.1.5	.632+ Bootstrap	6
2.2	Algorithms	6
3	Analysis	7
3.1	Simulated Data	8
3.2	Lymphoma and Lung Gene Microarray Data	8
3.3	Ovarian Proteomic Data	8
4	Results	8
5	Discussion	18

List of Tables

1	<i>Simulation Study Prediction Error Estimates with 10 features.</i> The estimate $\hat{\theta}_n$ (col 4) and standard deviation (col 5) based on learning sample of size 40. The estimate $\tilde{\theta}_n$ (rows 1-4) and standard deviation based on the remaining 260 observations. Bias (col 6) and MSE (col 7) reported for each resampling technique (col 1) and algorithm (col 3). The ten features with largest t -statistics used in algorithms. Minimums in bold	11
2	<i>Lymphoma Study Prediction Error Estimates with 10 genes.</i> Comparison of resampling method's MSE, bias, and standard deviation. Results shown are for the DDA algorithm using the top 10 genes as ranked by corresponding t -tests.	11
3	<i>Ovarian Study Prediction Error Estimates with 10 genes.</i> Comparison of resampling method's MSE, bias, and standard deviation. Results shown are for the DDA algorithm using the top 10 genes as ranked by corresponding t -tests.	12
4	<i>Simulation Study Prediction Error Estimates with 20 features.</i> The estimate $\hat{\theta}_n$ (col 4) and standard deviation (col 5) based on learning sample of size 40. The estimate $\tilde{\theta}_n$ (rows 1-4) and standard deviation based on the remaining 260 observations. Bias (col 6) and MSE (col 7) reported for each resampling technique (col 1) and algorithm (col 3). The 20 features with largest t -statistics used in algorithms. Minimums in bold	12
5	<i>Simulation Study MCCV Prediction Error Estimates with 10 features.</i> The estimate $\hat{\theta}_n$ (col 4) and standard deviation (col 5) based on learning sample of size 40. Bias (col 6) and MSE (col 7) reported for each MCCV fold (col 1) and algorithm (col 3). The ten features with largest t -statistics used in algorithms. Minimums in bold	13
6	<i>Repeated v-fold Resampling in the Simulation Study with 10 features.</i> To assess the effect of repeated resampling on v -fold CV a single v -fold CV is compared to the average of 10 repeats in the simulation study. The bias (col 7) and MSE (col 8) are reported over sample sizes of 40 and 80 (col 1), 2-,5-, and 10-fold CV (col 2), and four algorithms (col 4). The ten features with largest t -statistics are used in algorithms.	14

7	<i>Repeated Split-Sample Resampling in the Simulation Study with 10 features.</i> To assess the effect of repeated resampling on Split-Sample a single split is compared to the average of 10 splits in the simulation study. The bias (col 6) and MSE (col 7) are reported over sample sizes of 40, 80, and 120 (col 1) and four algorithms (col 3). The ten features with largest t -statistics are used in algorithms.	15
8	<i>Simulation Study Prediction Error Estimates without feature selection.</i> To assess the effect of no feature selection on resampling methods estimation, only 10 features were simulated and all 10 used in estimation. Results based on learning sample of 40 and a test sample of size 260. Absolute minimums in bold	16
9	<i>Resampling with and without Replacement.</i> The leave-one-out bootstrap and 3-fold MCCV estimate (col 3), standard deviation (col 4), bias (col 5), and MSE (col 6), over 3 samples sizes and 4 algorithms. Feature selection was used to select the top 10 ranked features by t -tests.	16

List of Figures

1	<i>Lung Study Prediction Error Estimates with DDA using 10 genes.</i>	17
---	---	----

1 Introduction

In genomic experiments one frequently encounters high dimensional data and small sample sizes. Microarrays simultaneously monitor expression levels for several thousands of genes. Proteomic profiling studies using SELDI-TOF (surface-enhanced laser desorption and ionization time-of-flight) measure size and charge of proteins and protein fragments by mass spectroscopy, and result in up to 15,000 intensity levels at prespecified mass values for each spectrum. Sample sizes in such experiments are typically less than 100.

In many studies observations are known to belong to predetermined classes and the task is to build predictors or classifiers for new observations whose class is unknown. Deciding which genes or proteomic measurements to include in the prediction is called *feature selection* and is a crucial step in developing a class predictor. Including too many noisy variables reduces accuracy of the prediction and may lead to over-fitting of data, resulting in promising but often non-reproducible results (Ransohoff, 2004).

Another difficulty is model selection with numerous classification models available. An important step in reporting results is assessing the chosen model's error rate, or generalizability. In the absence of independent validation data, a common approach to estimating predictive accuracy is based on some form of resampling the original data, e.g., cross-validation. These techniques divide the data into a learning set and a test set and range in complexity from the popular learning-test split to v -fold cross-validation, Monte-Carlo v -fold cross-validation, and bootstrap resampling. Few comparisons of standard resampling methods have been performed to date, and all of them exhibit limitations that make their conclusions inapplicable to most genomic settings. Early comparisons of resampling techniques in the literature are focussed on model selection as opposed to prediction error estimation (Breiman and Spector, 1992; Burman, 1989). In two recent assessments of resampling techniques for error estimation (Braga-Neto and Dougherty, 2004; Efron, 2004), feature selection was not included as part of the resampling procedures, causing the conclusions to be inapplicable to the high-dimensional setting.

We have performed an extensive comparison of resampling methods to estimate prediction error using simulated (large signal to noise ratio), microarray (intermediate signal to noise ratio) and proteomic data (low signal to noise ratio), encompassing increasing sample sizes with large numbers of features. The impact of feature selection on the performance of various cross validation methods is highlighted. The results elucidate the 'best' resampling techniques for future research involving high dimensional data to avoid overly optimistic assessment of the performance of a model.

2 Methods

In the prediction problem, one observes n independent and identically distributed (i.i.d.) random variables O_1, \dots, O_n with unknown distribution P . Each observation in O consists of an outcome Y with range \mathcal{Y} and a l -vector of measured covariates, or features, X with range \mathcal{X} , such that $O_i = (X_i, Y_i)$, $i = 1, \dots, n$. In microarray experiments X includes gene expression measurements, while in proteomic data, it includes the intensities at the mass over charge (m/z) values. X may also contain covariates such as a patient's age and/or histopathologic measurements. The outcome Y may be a continuous measure such as months to disease or a categorical measure such as disease status.

The goal in class prediction is to build a rule implementing the information from X in order to predict Y . The intention is that by building this rule based on the observations O_1, \dots, O_n a future unobserved outcome Y_0 can be predicted based on its corresponding measured features X_0 . If the outcome is continuous, then the rule, or predictor, ψ is defined as a mapping from the feature space \mathcal{X} onto the real line, i.e. $\psi : \mathcal{X} \rightarrow \mathbb{R}$. Consequently, $\hat{y} = \psi(x)$ denotes the predicted outcome based on the observed X . Such predictors can be built via regression (linear and non) or recursive binary partitioning such as Classification and Regression Trees (CART) (Breiman *et al.*, 1984). If the outcome Y is categorical it assumes one of K values. In this case, the rule ψ partitions the feature space \mathcal{X} into K disjoint and exhaustive groups G_k , where $k = 1, \dots, K$, such that $\hat{y} = k$ if $x \in G_k$. Standard statistical analyses include linear discriminant analysis (LDA) and diagonal discriminant classifiers (DDA), nearest neighbors (NN), and CART, as well as aggregate classifiers. Thorough discussions of the prediction problem and available algorithms can be found in Breiman *et al.* (1984); McLachlan (1992); Ripley (1996); Hastie *et al.* (2003).

The rule ψ can be written as $\psi(\cdot | P_n)$, where P_n denotes the empirical distribution of O and reflects the dependence of the built rule on the observed data. Loss functions may be employed to quantify the performance of a given rule. A common loss function for a continuous outcome Y is the squared error loss, $L(Y, \psi) = [Y - \psi(X)]^2$. With a categorical outcome Y , a popular choice is the indicator loss function, $L(Y, \psi) = I[Y \neq \psi(X)]$. A loss function could also incorporate differential misclassification costs (Breiman *et al.*, 1984).

For either type of outcome, the expected loss, or *risk*, is defined as:

$$\tilde{\theta} = R(\psi, P) = E_P[L(Y, \psi)] = \int L(y, \psi(x))dP(x, y). \quad (1)$$

The rule in (1) is constructed and evaluated upon the distribution P , as such, $\tilde{\theta}$ is referred to as the *asymptotic risk*. However, in reality P is unknown, thus, the rule based upon the observations O_1, \dots, O_n has an expected loss, or *conditional risk* (also known as the generalization error), defined as:

$$\tilde{\theta}_n = R(\psi(\cdot | P_n), P) = \int L(y, \psi(x | P_n))dP(x, y). \quad (2)$$

There are two impetuses for evaluating the conditional risk: model selection and performance assessment. In model selection, the goal is to find the one which minimizes the conditional risk over a collection of potential models. In performance assessment, the goal is to estimate the generalization error for a given model, i.e., assess how well it predicts the outcome of an observation not included in O .

In an ideal setting an independent data set would be available for the purposes of model selection and estimating the generalization error. Typically, however, one must use the observed sample O for model building, selection, and performance assessment. The simplest method for estimating the conditional risk is with the *resubstitution* or *apparent* error:

$$\hat{\theta}_n^{RS} = R(\psi(\cdot | P_n), P_n) = \int L(y, \psi(x | P_n))dP_n(x, y). \quad (3)$$

Here all n observations are used for constructing, selecting, and, subsequently, evaluating the prediction error of ψ . Consequently, the resubstitution risk estimate tends to underestimate the generalization error (Efron, 1983; McLachlan, 1992). To alleviate this biased estimation, resampling methods such as cross-validation or bootstrapping can be employed. In the next section, we describe these techniques and their implications in the framework of prediction error.

2.1 Resampling Methods

In the absence of a large, independent test set, there are numerous techniques for assessing prediction error by implementing some form of partitioning or resampling of the original observed data O . Each of these techniques involves dividing the data into a *learning set* and a *test set*. For purposes of model selection the learning set may further be divided into a *training set* and a *validation set*. We will focus solely on the partitioning of the data into learning and test sets for the purpose of estimating the generalization error.

To enhance a general discussion of resampling methods we define a binary random n -vector, $S_n \in \{0, 1\}^n$, which splits the observations into the desired subsets (Molinaro *et al.*, 2004). A realization of $S_n = (S_{n,1}, \dots, S_{n,n})$ prescribes a particular split of the entire data set of n observations into a learning set, $\{i \in \{1, \dots, n\} : S_{n,i} = 0\}$, and a test set, $\{i \in \{1, \dots, n\} : S_{n,i} = 1\}$. Let p be the proportion of observations in the test set. The empirical distributions of the learning and test sets are denoted by P_{n,S_n}^0 and P_{n,S_n}^1 , respectively. Importantly, S_n is independent of the empirical distribution of the complete data set of n observations P_n and the particular distribution of S_n defines the type of resampling method. Given S_n , the performance of any given estimator $\psi(\cdot | P_n)$ can be assessed via the *resampling conditional risk* estimate

$$\hat{\theta}_{n(1-p)} = E_{S_n} \int L(o, \psi(\cdot | P_{n,S_n}^0)) dP_{n,S_n}^1(o), \quad (4)$$

where S_n refers to binary split vectors for the entire data set of n observations and $p = \sum_i S_{i,n}/n$ is the proportion of n observations in the test set.

There are several considerations when selecting a resampling method. The first is sample size n . For v -fold cross-validation and bootstrap, Dudoit and van der Laan (2003) have shown that as $n \rightarrow \infty$ (and consequently $np \rightarrow \infty$) asymptotic optimality is achieved. However, no such results exist for finite samples. Other considerations are on the proportion p of the observations for the test set and the number of times the estimate is calculated. We address these considerations in the following sections and refer the reader to more detailed discussions in McLachlan (1992) and Davison and Hinkley (1997).

2.1.1 Split Sample

This popular resampling method, also known as the *learning-test split* or *holdout method* (McLachlan, 1992), entails a single partition of the data into a learning set and a test set based on a predetermined p . For example, $p = 1/3$ allots two-thirds of the data to the learning set and one third to the test set. The distribution of S_n places mass $1/2$ on two binary vectors which assign the n observations to the learning and test sets. The advantage of this method is the ease of computation. Also, since the classifier is developed only once, a completely specified algorithm for classifier development need not be available; the development can be more informal and subjective. There are two potential sources of bias inherent in this method: bias introduced by each individual observation contributing only to the learning or test set; and, bias due to a small learning set whereas both features and classifiers selected depend solely on the learning set. Because the learning set is smaller than the full data set, the test set error for a model built on the training set will tend to over-estimate the unknown generalization error for a model built on the full dataset.

2.1.2 v -fold Cross-Validation

This method randomly assigns the n observations to one of v partitions such that the partitions are near equal size. Subsequently, the learning set contains all but one of the partitions which is labeled the test set. The generalization error is assessed for each of the v test sets and then averaged over v . In this method, the distribution of S_n puts mass $1/v$ on the v binary vectors which assign each of the n observations to one of the v partitions. The proportion p is approximately equal to $1/v$. Both p and the number of averages can adversely or positively affect this estimate of error. For example, a larger v (e.g., $v = 10$) results in a smaller proportion p in the test set; thus, a higher proportion in the learning set decreasing the bias. In addition, the number of averages is equivalent to v and thus, may additionally decrease the bias.

2.1.3 Leave-One-Out Cross-Validation (LOOCV)

This is the most extreme case of v -fold cross-validation. In this method each observation is individually assigned to the test set, i.e., $v = n$ and $p = 1/n$ (Lachenbruch and Mickey, 1968; Geisser, 1975; Stone, 1974, 1977). The distribution of S_n places mass $1/n$ on the n binary vectors which assign each of the n observations to the learning and test sets. LOOCV and the corresponding $p = 1/n$ represent the best example of a bias-variance trade-off. It tends toward a small bias with elevated variance. In model selection, LOOCV has performed poorly compared to 10-fold cross-validation (Breiman and Spector, 1992). Due to the computational burden, LOOCV has not been a favored method for large samples and its behavior in estimating generalization error has not been thoroughly studied.

2.1.4 Monte Carlo Cross-Validation (MCCV)

MCCV randomly splits the sample into a learning and test set numerous times (e.g., 20, 50, or 1000 iterations). For each split $np = n(1/v)$ of the observations are labeled as the test set and $n(1-p) = n(1-1/v)$ as the learning set. For example, in MCCV with $v = 10$ each of 50 iterations allot 10% of the data to the test set and 90% to the learning set. The generalization error is assessed for each of the 50 test sets and subsequently averaged over the 50 iterations. The distribution of S_n puts mass $1/\binom{n}{np}$ on each of the binary vectors representing one split into a learning and test set. As the number of iterations increase the computational burden of MCCV is quite large. However, unless the iterations of random splits approaches infinity, the chance that each observation is included in a learning set and a test set (over all iterations) is small introducing a similar bias to that of the split sample approach (i.e., when each observation is either in the learning set or test set).

2.1.5 .632+ Bootstrap

Several variations of the bootstrap have been introduced to estimate the generalization error. The leave-one-out bootstrap ($\hat{\theta}_n^{BS}$) is based on a random sample drawn with replacement from n observations (Efron, 1983; Efron and Tibshirani, 1993). For each draw the observations left out (approximately $.368n$) serve as the test set. The learning set has approximately $.632n$ unique observations which leads to an overestimation of the prediction error (i.e., a decrease in the learning set leads to an increase in the bias). To correct for this two estimators have been suggested: the .632 bootstrap and the .632+ estimator. Both correct by adding the underestimated resubstitution error $\hat{\theta}_n^{RS}, \omega\hat{\theta}_n^{BS} + (1-\omega)\hat{\theta}_n^{RS}$. For the .632 bootstrap the weight ω is constant ($\omega = .632$), whereas for the .632+ bootstrap ω is determined based on the "no-information error rate" (Efron and Tibshirani, 1997). We focus on the latter as it is the most widely used in the literature and the most robust across different algorithms (Efron and Tibshirani, 1997).

2.2 Algorithms

Predictions of outcomes based on the observed X can employ parametric or non-parametric algorithms. If the outcome is continuous, predictors can be built using regression models or recursive binary partitioning like CART. If the outcome is categorical, algorithms which partition the feature space X into disjoint and exhaustive groups are used. In this manuscript, we limit our discussion to the classification of binary outcomes, i.e., $Y = 0$ or $Y = 1$, and thus, evaluate methods for the estimation of prediction error in the context of the following classification algorithms.

We calculate the LDA with the `lda` function in the `MASS` library of the statistical package `R` (Venables and Ripley, 1994; Ihaka and Gentleman, 1996). We use the function `dlda` in the library `supclust` in `R` to implement DDA (Dettling and Maechler, 2004). The library `supclust` also houses the function `nnr` for NN. CART classification is obtained using the library and function `rpart` in `R` (Breiman *et al.*, 1984; Therneau and Atkinson, 1997)

3 Analysis

The goal of this analysis is to ascertain differences between resampling methods in the estimation of the generalization error (presently, limited to the classification problem) in the presence of feature selection. We evaluate the influence of sample size, parametric to non-parametric classification methods, and large feature spaces on each resampling method’s ability to estimate the resampling conditional risk $\hat{\theta}_{n(1-p)}$ (Eq. 4) compared to that of the “true” conditional risk $\tilde{\theta}_n$ (Eq. 2). As such, a range of sample sizes ($n = 40$, $n = 80$, and $n = 120$), classification algorithms (LDA, DDA, NN, and CART), and data sets (simulated, microarray, and proteomic; see Sections 3.1–3.3) are utilized. Prior to discussing results, the general strategy for estimating the risks is explained followed by the specifics of each data set.

Each data set consists of N observations with N_1 cases and N_0 controls and l measured features. For $r = 1, \dots, R$ repetitions, a random sample of size n stratified by case/control status is selected from N such that the number of cases in the subsample ($n/2$) equals the number of controls. The stratification allows for equal representation of both cases and controls such that classification algorithms relying on majority consensus are not biased toward either (Quackenbush, 2004). This random sample, or subsample, plays two roles. First, it serves as a sample from which the resampling conditional risk $\hat{\theta}_{n(1-p)}$ can be estimated. This is accomplished by splitting the subsample into a learning and test set corresponding to each of the resampling methods. For each r , an estimate of $\hat{\theta}_{n(1-p)}$ is obtained for each resampling method with all four algorithms. In reality the distribution P of the observed data O is unknown and thus, so is the “true” conditional risk. In order to estimate $\tilde{\theta}_n$ in Eq. 2 we will use the complete observed data. As such, the subsample’s second role is to serve as the learning set and the remaining $N - n$ observations as the test set for an approximation of the conditional risk $\tilde{\theta}_n$.

To compare the resampling methods in Section 2.1, conditional risk estimates for each method are calculated and compared to each other and the truth (i.e., the conditional risk). This evaluation is based on the mean squared error (MSE) and bias, calculated as follows:

$$MSE = \frac{1}{r} \sum_{r=1}^R (\hat{\theta}_{n,r} - \tilde{\theta}_{n,r})^2$$

$$Bias = \frac{1}{r} \sum_{r=1}^R (\hat{\theta}_{n,r} - \tilde{\theta}_{n,r}),$$

where $\hat{\theta}_{n,r}$ is the resampling conditional risk and $\tilde{\theta}_{n,r}$ is the conditional risk for the r th repetition. In all results the total number of repetitions is set at 100, i.e., $R = 100$.

There were several attempts to examine the effect of varying p on those resampling methods which allow user-defined test set proportions (i.e., v -fold cross-validation, MCCV, and split sample). For v -fold cross validation, 2, 5, and 10-fold were explored. In MCCV, both p and the number of MCCV repetitions affect the estimation, thus, test set proportions of $p = 0.5$, $p = 0.2$, $p = 0.1$ as well as repetitions of 20, 50, and 1000 were run. In split sample estimation test set proportions of both $p = 1/3$ and $p = 1/2$ were examined to assess the bias/variance trade-off.

Given the high-dimensional structure of each data set (i.e., large l), feature selection is an important step before running any of the algorithms. Feature selection must be based on the learning set within each resampling, otherwise additional bias is introduced (Simon *et al.*, 2003). This correct approach to feature selection within cross-validation has been referred to as *honest* or *complete* (Quackenbush, 2004). There are many methods available for feature selection, here univariate t -tests are used. Initially components of X with the largest 10 absolute value t -test statistics are considered. Subsequently, the largest 20 are discussed.

All simulations and analyses were implemented in R (Ihaka and Gentleman, 1996).

3.1 Simulated Data

The simulated data sets are generated as described in [Bura and Pfeiffer \(2003\)](#). Each data set contains $N = 300$ observations with 750 covariates, representing patients and features (or genes), respectively. Half of the observations (i.e., 150) are labeled controls ($Y = 0$) and half cases ($Y = 1$). Of the 750 genes, 8 are associated with disease and the others are non-predictive. The controls are simulated from a multivariate normal distribution with a mean of 0 and covariance matrix Σ . The cases have 98% non-differentially expressed genes which are generated from the same $N(0, \Sigma)$ as the controls. The 2% of the genes that are differentially expressed are generated from a mixture of two multivariate normals with means μ_1 and μ_2 and covariance structure Σ . The mixing probability is 0.5. The covariance matrix $\Sigma = (\sigma_{ij})$ has a block structure with $\sigma_{ij} = 0.2$ for $|j - i| \leq 5$ and zero otherwise. Estimates of $\hat{\theta}_{n(1-p)}$ and $\tilde{\theta}_n$ are based on learning samples of size 40, 80, and 120 and test sets of size 260, 220, and 180, respectively.

3.2 Lymphoma and Lung Gene Microarray Data

Lymphoma Cancer Data The lymphoma microarray data set is publicly available and focuses on diffuse large-B-cell lymphoma ([Rosenwald et al., 2002](#)). In this study there are 7399 genes on the microarray and 240 patients. For the purposes of this analysis, the outcome variable represents the lymphoma subtype: activated B-cell for $Y = 0$ and germinal-center B-cell for $Y = 1$. This is an example of a moderate signal to noise ratio data set as the subgroups do not separate perfectly based on the microarray observations ([Wright et al., 2003](#)). Estimates of $\hat{\theta}_{n(1-p)}$ and $\tilde{\theta}_n$ are based on learning samples of size 40, 80, and 120 and test sets of size 200, 160, and 120, respectively.

Lung Cancer Data The second study uses oligonucleotide microarrays to measure 12601 transcript sequences for 186 lung tumor samples ([Bhattacharjee et al., 2001](#)). For our analysis, the outcome represents the 139 adenocarcinomas as $Y = 0$ and the remaining 47 tumors as $Y = 1$. Estimates of $\hat{\theta}_{n(1-p)}$ and $\tilde{\theta}_n$ are based on learning samples of size 40 and 80 and test sets of size 146 and 106, respectively. This is another example of a moderate signal to noise ratio data set.

3.3 Ovarian Proteomic Data

The proteomic data set consists of 164 SELDI-TOF measurements from NCI/Mayo Clinic serum samples. These data are part of a study designed to validate previously identified proteomic markers for ovarian cancer. The readings are from fraction 4, IMAC30 ProteinChip arrays, read at high and low energy settings in a PCS4000 ProteinChip Reader (Ciphergen Biosystems, Inc., Fremont, CA). The spectra were externally calibrated for mass, internally normalized for intensity using total ion current, and baseline subtracted. Peaks were manually selected and the intensity recorded.

Of the $n = 164$ observations, 45 are ovarian cancer cases and 119 controls. Estimates of $\hat{\theta}_{n(1-p)}$ and $\tilde{\theta}_n$ are based on learning samples of size 40 and 80 and test sets of size 144 and 104, respectively. Given the nature of proteomic data as well as the naive algorithms implemented this will serve as a low signal to noise example.

4 Results

A limited number of tables and figures are displayed in this manuscript; however, all tables and figures can be found at <http://linus.nci.nih.gov/brb/TechReport.htm>.

Simulation Study Results. For $n = 40$, LOOCV and 10-fold CV have the smallest MSE and bias followed by 5-fold CV and then .632+ (Table 1). The largest MSE and bias occur with 2-fold CV and Split Sample with $p = 1/2$. For $n = 80$ and $n = 120$ the differences among these methods diminish. For $n = 40$ and

$n = 80$, .632+ has the smallest standard deviation followed by 10-fold CV, LOOCV, and 5-fold CV. The only exception is for LDA and NN at $n = 80$, when LOOCV and 10-fold CV have the smallest. At $n = 120$, the differences among these methods diminish.

Lymphoma and Lung Study Results. In the lymphoma study, for $n = 40, 80$, and 120, .632+, LOOCV, 5- and 10-fold CV have the smallest MSE and bias. The two split samples and 2-fold CV have the largest MSE and bias. Similar to the simulation study, .632+ has the smallest standard deviation across the algorithms and sample sizes, while both Split Samples do by far the worst. The results with the DDA algorithm are shown in Table 2. In the lung tumor study, for $n = 40$ and $n = 80$, .632+, LOOCV, 5- and 10-fold CV have the smallest MSE and bias. The two split samples and 2-fold CV have the largest MSE and bias. Similar to the simulation and lymphoma study, .632+ has the smallest standard deviation across the algorithms and sample sizes. The results for the lung tumor study with the DDA algorithm are shown in Figure 1.

Ovarian Study Results. For $n = 40$ to $n = 80$, LOOCV and .632+ have the smallest MSE followed by 5- and 10-fold CV. As for bias, 10-fold CV, .632+, and LOOCV vie for the smallest. The largest MSE and bias occurs with the Split Samples and 2-fold CV. Again .632+ has the smallest standard deviation across algorithms and sample sizes; however, the discrepancy is much smaller than in the other two studies. The Split Samples have the largest standard deviations. The results with the DDA algorithm are shown in Table 3.

Twenty features. All analyses were repeated with selecting the 20 features having the largest t -test statistics. The ranking of the resampling methods remained the same. Results from choosing the top 20 genes in the simulation study for $n = 40$ are shown in Table 4.

MCCV. The MCCV estimates show a slight decrease in variance in comparison to those of v -fold CV. Interestingly, the advantage of increasing the MCCV iterations from 20 to 50 to 1000 is minimal, i.e., the MSE is equivalent over the iterations while the bias and variance slightly decrease. The MCCV results for the simulation study at $n=40$ are shown in Table 5.

Repeated Resampling. We examined the effect of repeated resampling on 2-, 5-, and 10-fold CV and Split Sample with $p = 1/3$ for the three sample sizes and four algorithms. Each was repeated 10 and 30 times. Interestingly, there was minimal improvement when increased from 10 to 30 repeats. However when increasing repeats from 1 to 10 (or 30), all standard deviations decreased (up to 50%). The MSE either decreased (up to 35%) or stayed similar, which was also true for the bias except in Split Sample for $n=40$ and 2-fold CV for $n=40$ and $n=80$. The results for repeated v -fold CV are shown in Table 6 and those for repeated Split-Sample in Table 7.

Dimensionality of Feature Space. In the simulations of Efron and Tibshirani (1997), .632+ outperformed LOOCV and 10-fold CV. For example, in their experiment 22, with 10 variables and 36 patients, the MSE was .040 for .632+ and .058 for LOOCV. However, in our simulations with $n=40$ (Table 1) .632+ does not fair so well, particularly with regard to bias. To investigate the differences between our simulations and those in Efron and Tibshirani we decreased the dimensions of the feature space to a total of 10 variables instead of 750. The results are shown in Table 8 for the sample size of 40. With low dimension the large bias of the bootstrap is substantially reduced and the .632+ does as well or better than LOOCV and 10-fold CV.

Resampling with and without Replacement. To understand the ramification of resampling with replacement as it pertains to the bootstrap estimates we compared the leave-one-out bootstrap estimate (Section 2.1.5) to the 3-fold MCCV. The 3-fold MCCV randomly selects $.666n$ for the learning set and $.333n$ for the test set. This is repeated numerous times and the estimates averaged. Therefore the 3-fold MCCV is equivalent to the leave-one-out bootstrap except it employs resampling without replacement. Table 9 displays the Simulation Study results for the two estimates using 50 iterations for both. Interestingly, the bias and MSE for the leave-one-out bootstrap are roughly double that of 3-fold MCCV. The only two distinct

differences between the two methods are the replicate copies in the learning set inherent in the bootstrap estimate and the fact that **on average** $.632n$ unique observations are in the learning sample for the leave-one-out bootstrap whereas there are *always* $.666n$ in the learning sample for the 3-fold MCCV. Both of these factors may contribute to the increase in bias and MSE. The results for $n = 40$ can be compared to the resubstitution and $.632+$ estimates for the same sample size in Table 1.

Table 1: *Simulation Study Prediction Error Estimates with 10 features.* The estimate $\hat{\theta}_n$ (col 4) and standard deviation (col 5) based on learning sample of size 40. The estimate $\tilde{\theta}_n$ (rows 1-4) and standard deviation based on the remaining 260 observations. Bias (col 6) and MSE (col 7) reported for each resampling technique (col 1) and algorithm (col 3). The ten features with largest t -statistics used in algorithms. Minimums in **bold**.

Estimator	p	Algorithm	Est	St.Dev	Bias	MSE
$\tilde{\theta}_n$	0.87	LDA	0.078	0.093		
		DDA	0.160	0.086		
		NN	0.042	0.084		
		CART	0.121	0.133		
$\hat{\theta}_n^{RS}$	1	LDA	0.007	0.016	-0.071	0.012
		DDA	0.04	0.039	-0.12	0.019
		NN	0	0	-0.042	0.009
		CART	0.039	0.052	-0.082	0.015
v -fold CV	0.5	LDA	0.357	0.126	0.279	0.097
		DDA	0.342	0.106	0.182	0.052
		NN	0.277	0.135	0.235	0.077
		CART	0.430	0.121	0.309	0.134
	0.2	LDA	0.161	0.127	0.083	0.017
		DDA	0.208	0.086	0.048	0.012
		NN	0.108	0.102	0.066	0.011
		CART	0.284	0.117	0.163	0.055
	0.1	LDA	0.118	0.120	0.040	0.008
		DDA	0.177	0.087	0.017	0.007
		NN	0.078	0.102	0.036	0.005
		CART	0.189	0.104	0.068	0.024
LOOCV	0.025	LDA	0.092	0.115	0.014	0.008
		DDA	0.164	0.096	0.004	0.007
		NN	0.058	0.103	0.016	0.005
		CART	0.146	0.125	0.025	0.018
SPLIT	0.333	LDA	0.205	0.184	0.127	0.053
		DDA	0.243	0.138	0.083	0.034
		NN	0.145	0.169	0.103	0.044
		CART	0.371	0.174	0.25	0.121
	0.5	LDA	0.348	0.185	0.270	0.113
		DDA	0.344	0.139	0.184	0.062
		NN	0.265	0.177	0.223	0.086
		CART	0.438	0.155	0.317	0.147
.632+ 50 reps	\approx .368	LDA	0.274	0.084	0.196	0.047
		DDA	0.286	0.074	0.126	0.028
		NN	0.200	0.070	0.158	0.032
		CART	0.387	0.080	0.266	0.100

Table 2: *Lymphoma Study Prediction Error Estimates with 10 genes.* Comparison of resampling method's MSE, bias, and standard deviation. Results shown are for the DDA algorithm using the top 10 genes as ranked by corresponding t -tests.

Resampling Method	$n = 40$			$n = 80$			$n = 120$		
	St.Dev	Bias	MSE	St.Dev	Bias	MSE	St.Dev	Bias	MSE
2-fold CV	0.085	0.038	0.01	0.043	0.002	0.004	0.031	0.0	0.003
5-fold CV	0.07	0.004	0.007	0.045	-0.008	0.005	0.032	-0.006	0.003
10-fold CV	0.063	-0.007	0.006	0.036	-0.009	0.003	0.031	-0.006	0.003
LOOCV	0.072	-0.019	0.008	0.04	-0.013	0.004	0.033	-0.004	0.003
SPLIT 1/3	0.119	0.001	0.017	0.071	0.0	0.007	0.059	-0.004	0.005
SPLIT 1/2	0.117	0.037	0.018	0.058	0.001	0.005	0.046	-0.001	0.004
.632+	0.049	-0.006	0.004	0.025	-0.02	0.003	0.018	-0.015	0.002

Table 3: *Ovarian Study Prediction Error Estimates with 10 genes.* Comparison of resampling method’s MSE, bias, and standard deviation. Results shown are for the DDA algorithm using the top 10 genes as ranked by corresponding t -tests.

Resampling Method	$n = 40$			$n = 80$		
	St.Dev	Bias	MSE	St.Dev	Bias	MSE
2-fold CV	0.098	0.026	0.015	0.05	0.004	0.007
5-fold CV	0.082	0.0	0.012	0.039	-0.005	0.006
10-fold CV	0.082	-0.01	0.011	0.036	-0.005	0.005
LOOCV	0.079	-0.004	0.011	0.037	-0.004	0.006
SPLIT 1/3	0.133	-0.002	0.022	0.075	-0.009	0.009
SPLIT 1/2	0.113	0.027	0.018	0.071	0.013	0.01
.632+	0.075	-0.006	0.011	0.028	-0.014	0.005

Table 4: *Simulation Study Prediction Error Estimates with 20 features.* The estimate $\hat{\theta}_n$ (col 4) and standard deviation (col 5) based on learning sample of size 40. The estimate $\tilde{\theta}_n$ (rows 1-4) and standard deviation based on the remaining 260 observations. Bias (col 6) and MSE (col 7) reported for each resampling technique (col 1) and algorithm (col 3). The 20 features with largest t -statistics used in algorithms. Minimums in **bold**.

Estimator	p	Algorithm	Est	St.Dev	Bias	MSE
$\tilde{\theta}_n$	0.87	LDA	0.117	0.086		
		DDA	0.215	0.06		
		NN	0.043	0.078		
		CART	0.123	0.135		
v -fold CV	0.5	LDA	0.37	0.109	0.253	0.078
		DDA	0.343	0.087	0.128	0.026
		NN	0.191	0.125	0.148	0.036
		CART	0.433	0.115	0.31	0.134
v -fold CV	0.2	LDA	0.182	0.109	0.065	0.014
		DDA	0.25	0.089	0.035	0.012
		NN	0.072	0.093	0.029	0.006
		CART	0.277	0.105	0.154	0.049
v -fold CV	0.1	LDA	0.147	0.111	0.03	0.008
		DDA	0.232	0.09	0.017	0.009
		NN	0.06	0.088	0.017	0.003
		CART	0.195	0.11	0.072	0.024
LOOCV	0.025	LDA	0.138	0.125	0.021	0.008
		DDA	0.234	0.106	0.019	0.012
		NN	0.049	0.087	0.006	0.002
		CART	0.144	0.122	0.021	0.011
SPLIT	0.333	LDA	0.262	0.172	0.145	0.051
		DDA	0.288	0.134	0.073	0.028
		NN	0.098	0.148	0.055	0.022
		CART	0.377	0.174	0.254	0.124
v -fold CV	0.5	LDA	0.362	0.156	0.245	0.084
		DDA	0.341	0.117	0.126	0.031
		NN	0.177	0.17	0.134	0.047
		CART	0.442	0.149	0.319	0.144
.632+ 50 reps	\approx .368	LDA	0.301	0.084	0.184	0.04
		DDA	0.311	0.074	0.096	0.017
		NN	0.148	0.072	0.105	0.017
		CART	0.395	0.080	0.272	0.104

Table 5: *Simulation Study MCCV Prediction Error Estimates with 10 features.* The estimate $\hat{\theta}_n$ (col 4) and standard deviation (col 5) based on learning sample of size 40. Bias (col 6) and MSE (col 7) reported for each MCCV fold (col 1) and algorithm (col 3). The ten features with largest t -statistics used in algorithms. Minimums in **bold**.

Estimator	Iterations	Algorithm	Est	St.Dev	Bias	MSE
MCCV 2	20	LDA	0.373	0.067	0.295	0.095
		DDA	0.35	0.066	0.19	0.048
		NN	0.269	0.069	0.227	0.06
		CART	0.449	0.061	0.328	0.133
	50	LDA	0.371	0.065	0.293	0.093
		DDA	0.349	0.068	0.189	0.048
		NN	0.269	0.067	0.227	0.059
		CART	0.448	0.058	0.327	0.132
	1000	LDA	0.369	0.06	0.291	0.092
		DDA	0.349	0.066	0.189	0.048
		NN	0.272	0.064	0.23	0.061
		CART	0.445	0.054	0.324	0.13
MCCV 5	20	LDA	0.16	0.121	0.082	0.014
		DDA	0.214	0.081	0.054	0.011
		NN	0.103	0.1	0.061	0.009
		CART	0.266	0.088	0.145	0.043
	50	LDA	0.159	0.117	0.081	0.014
		DDA	0.215	0.076	0.055	0.01
		NN	0.105	0.097	0.063	0.009
		CART	0.273	0.082	0.152	0.044
	1000	LDA	0.159	0.112	0.081	0.014
		DDA	0.213	0.074	0.053	0.01
		NN	0.105	0.094	0.063	0.009
		CART	0.277	0.074	0.156	0.045
MCCV 10	20	LDA	0.123	0.122	0.045	0.008
		DDA	0.183	0.086	0.023	0.008
		NN	0.077	0.105	0.035	0.005
		CART	0.196	0.112	0.075	0.023
	50	LDA	0.121	0.116	0.043	0.008
		DDA	0.182	0.079	0.022	0.007
		NN	0.074	0.095	0.032	0.005
		CART	0.191	0.099	0.07	0.021
	1000	LDA	0.116	0.112	0.038	0.008
		DDA	0.18	0.078	0.02	0.006
		NN	0.074	0.094	0.032	0.005
		CART	0.191	0.092	0.07	0.02

Table 6: *Repeated v-fold Resampling in the Simulation Study with 10 features.* To assess the effect of repeated resampling on v -fold CV a single v -fold CV is compared to the average of 10 repeats in the simulation study. The bias (col 7) and MSE (col 8) are reported over sample sizes of 40 and 80 (col 1), 2-,5-, and 10-fold CV (col 2), and four algorithms (col 4). The ten features with largest t -statistics are used in algorithms.

n	v -fold	Repeats	Algorithm	Est	St.Dev	Bias	MSE	
40	2	1	LDA	0.357	0.126	0.279	0.097	
			DDA	0.342	0.106	0.182	0.052	
			NN	0.277	0.135	0.235	0.077	
	5	10	CART	0.43	0.121	0.309	0.134	
			LDA	0.37	0.065	0.292	0.094	
			DDA	0.348	0.07	0.188	0.048	
	80	2	1	NN	0.269	0.069	0.227	0.06
				CART	0.448	0.061	0.327	0.132
				LDA	0.161	0.127	0.083	0.017
		5	10	DDA	0.208	0.086	0.048	0.012
				NN	0.108	0.102	0.066	0.011
				CART	0.284	0.117	0.163	0.055
80		2	1	LDA	0.159	0.114	0.081	0.014
				DDA	0.209	0.073	0.049	0.009
				NN	0.104	0.095	0.062	0.009
		5	10	CART	0.278	0.08	0.157	0.046
				LDA	0.118	0.12	0.04	0.008
				DDA	0.177	0.087	0.017	0.007
	80	2	1	NN	0.078	0.102	0.036	0.005
				CART	0.189	0.104	0.068	0.024
				LDA	0.116	0.109	0.038	0.008
		5	10	DDA	0.179	0.081	0.019	0.006
				NN	0.075	0.094	0.033	0.005
				CART	0.192	0.095	0.071	0.021
80		2	1	LDA	0.074	0.078	0.047	0.007
				DDA	0.14	0.08	0.069	0.01
				NN	0.037	0.066	0.027	0.004
		5	10	CART	0.126	0.098	0.084	0.016
				LDA	0.08	0.068	0.053	0.006
				DDA	0.151	0.057	0.08	0.009
	80	2	1	NN	0.047	0.058	0.037	0.004
				CART	0.128	0.051	0.086	0.01
				LDA	0.033	0.046	0.006	0.001
		5	10	DDA	0.086	0.056	0.015	0.002
				NN	0.015	0.033	0.005	0.001
				CART	0.049	0.044	0.007	0.002
80		2	1	LDA	0.033	0.042	0.006	0.001
				DDA	0.084	0.054	0.013	0.002
				NN	0.015	0.032	0.005	0.001
		5	10	CART	0.048	0.037	0.006	0.002
				LDA	0.029	0.037	0.002	0.001
				DDA	0.076	0.053	0.005	0.001
	80	2	1	NN	0.013	0.031	0.003	0
				CART	0.045	0.042	0.003	0.002
				LDA	0.029	0.034	0.002	0.001
		5	10	DDA	0.076	0.052	0.005	0.001
				NN	0.012	0.026	0.002	0
				CART	0.044	0.036	0.002	0.002

Table 7: *Repeated Split-Sample Resampling in the Simulation Study with 10 features.* To assess the effect of repeated resampling on Split-Sample a single split is compared to the average of 10 splits in the simulation study. The bias (col 6) and MSE (col 7) are reported over sample sizes of 40, 80, and 120 (col 1) and four algorithms (col 3). The ten features with largest t -statistics are used in algorithms.

n	Repeats	Algorithm	Est	St.Dev	Bias	MSE
40	1	LDA	0.205	0.184	0.127	0.053
		DDA	0.243	0.138	0.083	0.034
		NN	0.145	0.169	0.103	0.044
		CART	0.371	0.174	0.25	0.121
	10	LDA	0.24	0.121	0.162	0.039
		DDA	0.271	0.088	0.111	0.025
		NN	0.171	0.107	0.129	0.029
		CART	0.37	0.101	0.249	0.096
	30	LDA	0.24	0.112	0.162	0.036
		DDA	0.267	0.08	0.107	0.022
		NN	0.17	0.1	0.128	0.025
		CART	0.364	0.092	0.243	0.091
80	1	LDA	0.037	0.052	0.01	0.002
		DDA	0.11	0.087	0.039	0.009
		NN	0.02	0.055	0.01	0.002
		CART	0.043	0.052	0.001	0.004
	10	LDA	0.041	0.052	0.014	0.002
		DDA	0.102	0.059	0.031	0.003
		NN	0.021	0.045	0.011	0.001
		CART	0.057	0.041	0.015	0.002
	30	LDA	0.043	0.053	0.016	0.002
		DDA	0.102	0.057	0.031	0.003
		NN	0.022	0.044	0.012	0.001
		CART	0.061	0.041	0.019	0.002
120	1	LDA	0.028	0.038	0.007	0.001
		DDA	0.068	0.064	0.016	0.003
		NN	0.012	0.03	0.003	0.001
		CART	0.034	0.039	0.008	0.002
	10	LDA	0.027	0.03	0.006	0
		DDA	0.069	0.044	0.017	0.001
		NN	0.012	0.024	0.003	0
		CART	0.036	0.026	0.01	0.001
	30	LDA	0.027	0.028	0.006	0
		DDA	0.071	0.044	0.019	0.001
		NN	0.012	0.024	0.003	0
		CART	0.037	0.023	0.011	0.001


Table 8: *Simulation Study Prediction Error Estimates without feature selection.* To assess the effect of no feature selection on resampling methods estimation, only 10 features were simulated and all 10 used in estimation. Results based on learning sample of 40 and a test sample of size 260. Absolute minimums in **bold**.

Estimator	p	Algorithm	Est	St.Dev	Bias	MSE	
$\tilde{\theta}_n$	0.87	LDA	0.026	0.028			
		DDA	0.073	0.058			
		NN	0.010	0.017			
		CART	0.099	0.092			
v -fold CV	0.5	LDA	0.067	0.060	0.041	0.005	
		DDA	0.106	0.079	0.033	0.009	
		NN	0.011	0.025	0.001	0	
		CART	0.304	0.088	0.205	0.063	
	0.2	LDA	0.034	0.045	0.008	0.002	
		DDA	0.085	0.049	0.012	0.003	
		NN	0.011	0.024	0.001	0	
		CART	0.158	0.072	0.059	0.012	
	0.1	LDA	0.032	0.041	0.006	0.001	
		DDA	0.074	0.048	0.001	0.002	
		NN	0.010	0.021	0	0	
		CART	0.118	0.063	0.019	0.006	
LOOCV	0.025	LDA	0.028	0.040	0.002	0.001	
		DDA	0.072	0.049	-0.001	0.002	
		NN	0.010	0.022	0	0	
		CART	0.110	0.075	0.011	0.006	
SPLIT	0.333	LDA	0.046	0.076	0.020	0.005	
		DDA	0.066	0.085	-0.007	0.008	
		NN	0.007	0.029	-0.003	0.001	
		CART	0.265	0.116	0.166	0.047	
	0.5	LDA	0.073	0.078	0.047	0.007	
		DDA	0.093	0.099	0.020	0.013	
		NN	0.010	0.028	0	0.001	
		CART	0.308	0.114	0.209	0.071	
	.632+ 50 reps	$\approx .368$	LDA	0.037	0.036	0.011	0.001
			DDA	0.085	0.036	0.012	0.003
			NN	0.008	0.016	-0.002	0
			CART	0.160	0.034	0.061	0.010


Table 9: *Resampling with and without Replacement.* The leave-one-out bootstrap and 3-fold MCCV estimate (col 3), standard deviation (col 4), bias (col 5), and MSE (col 6), over 3 samples sizes and 4 algorithms. Feature selection was used to select the top 10 ranked features by t -tests.

n	Alg	Leave-one-out Bootstrap				3-fold MCCV			
		Est	St.Dev(Est)	Bias	MSE	Est	St.Dev(Est)	Bias	MSE
n=40	LDA	0.331	0.075	0.252	0.072	0.242	0.101	0.164	0.035
	DDA	0.337	0.075	0.177	0.044	0.270	0.072	0.110	0.022
	NN	0.259	0.072	0.217	0.055	0.167	0.083	0.125	0.022
	CART	0.414	0.065	0.296	0.114	0.377	0.085	0.256	0.094
n=80	LDA	0.07	0.063	0.043	0.004	0.044	0.053	0.017	0.002
	DDA	0.146	0.058	0.074	0.008	0.104	0.058	0.033	0.003
	NN	0.046	0.056	0.036	0.003	0.022	0.043	0.012	0.001
	CART	0.098	0.047	0.057	0.006	0.062	0.039	0.020	0.002
n=120	LDA	0.032	0.033	0.011	0.001	0.026	0.026	0.005	0
	DDA	0.088	0.045	0.036	0.002	0.068	0.043	0.016	0.001
	NN	0.016	0.030	0.007	0	0.012	0.023	0.003	0
	CART	0.048	0.025	0.022	0.001	0.038	0.022	0.012	0.001

DDA and Resampling in Lung Cancer Study



DDA and Resampling in Lung Cancer Study



DDA and Resampling in Lung Cancer Study




Figure 1: Lung Study Prediction Error Estimates with DDA using 10 genes.

5 Discussion

Estimation of prediction error when confronted with a multitude of covariates and small sample sizes is a relatively new problem. Feature selection, sample size, and signal to noise ratio have important influences on the relative performance of resampling methods. We have evaluated resampling methods for use in high dimensional classification problems using a range of sample sizes, algorithms, and signals. Some general conclusions may be summarized as follows:

1. **With small sample sizes, the Split Sample method and 2-fold CV perform very poorly.** This poor performance is primarily due to a large positive bias resulting from use of a reduced training set size which severely impairs its ability to effectively select features and fit a model. The large bias contributes to a large MSE.
2. **LOOCV generally performs very well with regard to MSE and bias.** The only exception is when an unstable classifier (e.g., CART) is used in the presence of a weak signal. In this setting, the larger MSE is attributed to LOOCV's increased variance.
3. **10-fold CV prediction error estimates approximate those of LOOCV in almost all settings.** For computationally burdensome analyses, 10-fold CV may be preferable to LOOCV. Additionally, in the simulated data, repeated resamplings (the average of 10 repeats) reduce the MSE, bias, and variance of 10-fold CV.
4. **The .632+ prediction error estimate performs best with moderate to weak signal to noise ratios.** Previous studies have found the bootstrap variants superior to LOOCV and v -fold CV; however, these studies did not include feature selection. As seen in Table 1, honest resampling in small samples with strong signal suggest that LOOCV and 10-fold CV are in fact better than the .632+ bootstrap. This discrepancy fades when feature selection is discarded (Table 8) and when the signal decreases as seen in the Lymphoma and Ovarian data sets (Tables 2 and 3). Additional glimpses into the bootstrap estimate (Table 9) indicate that the sampling with replacement increases the MSE and bias substantially over 3-fold MCCV (i.e., resampling without replacement).
5. MCCV does not decrease the MSE or bias enough to warrant its use over v -fold CV.
6. As the sample size grows the differences among the resampling methods decrease. Additionally, as the signal decreases from strong in the simulated data to rather weak in the ovarian data the discrepancies between the methods diminish.

In future work we will compare the resampling methods for continuous outcomes and continue to explore the behavior of the bootstrap estimates. Additionally, the effect of feature selection method may play an important role in prediction and needs further investigation.

Acknowledgement

A.M.M. was supported by the Cancer Prevention Fellowship Program, DCP/NCI/NIH. The authors thank Mark J. van der Laan for fruitful discussions.

References

- Bhattacharjee, A., *et al.* (2001) Classification of human lung carcinomas by mRNA expression profiling reveals distinct adenocarcinoma subclasses. *PNAS*, **98**(24), 13790–13795.
- Braga-Neto, U.M., Dougherty, E.R. (2004) Is Cross-validation for Small-sample Microarray Classification? *Bioinformatics*, **20**(3), 374–380.

- Breiman, L., Friedman, J.H., Olshen, R.A., Stone, C.J. (1984) *Classification and Regression Trees*. Wadsworth and Brooks/Cole, Monterey, CA.
- Breiman, L., Spector, P. (1992) Submodel selection and evaluation in regression. The X-random case. *International Statistical Review*, **60**, 291-319.
- Bura, E., Pfeiffer, R.M. (2003) Graphical methods for class prediction using dimension reduction techniques on DNA microarray data. *Bioinformatics*, **19**(10), 1252-1258.
- Burman, P. (1989) A comparative study of ordinary cross-validation, v -fold cross-validation and the repeated learning-testing methods. *Biometrika*, **76**(3), 503-14.
- Davison, A.C., Hinkley, D.V. (1997) *Bootstrap methods and their application*. Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press, Cambridge, United Kingdom.
- Dettling, M., Maechler, M. (2004) Methodology for Supervised Grouping of Predictor Variables. <http://lib.stat.cmu.edu/R/CRAN/src/contrib/Descriptions/supclust.html>
- Dudoit, S., van der Laan, M.J. (2003) Asymptotics of Cross-Validated Risk Estimation in Model Selection and Performance Assessment. Technical Report 126, U.C. Berkeley Division of Biostatistics Working Paper Series. URL <http://www.bepress.com/ucbbiostat/paper126>
- Efron, B. (1983) Estimating the Error Rate of a Prediction Rule: Improvement on Cross-Validation. *Journal of the American Statistical Association*, **78**(382), 316-331.
- Efron, B., Tibshirani, R.J. (1993) *An Introduction to the Bootstrap*. Monographs on Statistics and Applied Probability 57. Chapman & Hall.
- Efron, B., Tibshirani, R.J. (1997) Improvements on Cross-Validation: The .632+ Bootstrap Method. *Journal of the American Statistical Association*, **92**(438), 548-560.
- Efron, B. (2004) The Estimation of Prediction Error: Covariance Penalties and Cross-Validation. *Journal of the American Statistical Association*, **99**(467), 619-642.
- Geisser, S. (1975) The Predictive Sample Reuse Method With Applications. *Journal of the American Statistical Association*, **70**, 320-328.
- Hastie, T., Tibshirani, R., Friedman, J. (2003) *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer Series in Statistics. Springer, 1st edition.
- Ihaka, R., Gentleman, R. (1996) R: A language for data analysis and graphics. *Journal of Computational and Graphical Statistics*, **5**, 299-314.
- Lachenbruch, P.A., Mickey, M.R. (1968) Estimation of error rates in discriminant analysis. *Technometrics*, **10**, 1-11.
- McLachlan, G.J. (1992) *Discriminant Analysis and Statistical Pattern Recognition*. John Wiley & Sons, Inc..
- Molinaro, A.M., Dudoit, S., van der Laan, M.J. (2004) Tree-based Multivariate Regression and Density Estimation with Right-censored Data. *Journal of Multivariate Analysis*, **90**(1):154-177.
- Quackenbush, J. (2004) Meeting the Challenges of Functional Genomics: From the Laboratory to the Clinic. *Preclinica***2**(5), 313-316.
- Ransohoff, D.F. (2004) Rules of evidence for cancer molecular marker discovery and validation. *Nature Reviews/Cancer* **4**, 309-313.
- Ripley, B.D. (1996) *Pattern Recognition and Neural Networks*. Cambridge University Press.

- Rosenwald, A., Wright, G., Chan, W.C., et al. (2002) The use of molecular profiling to predict survival after chemotherapy for diffuse large-B-cell lymphoma. *The New England Journal of Medicine*, **346**(25), 1937-1946.
- Simon, R., Radmacher, M.D., Dobbin, K., McShane, L.M. (2003) Pitfalls in the Use of DNA Microarray Data for Diagnostic and Prognostic Classification. *Journal of the National Cancer Institute*, **95**(1), 14-18.
- Stone, M. (1974) Cross-validated choice and assessment of statistical predictions. *Journal of the Royal Statistical Society, Series B*, **36**, 111-147.
- Stone, M. (1977) Asymptotics for and against cross-validation. *Biometrika*, **64**(1), 29-35.
- Therneau, T., Atkinson, E. (1997) An Introduction to Recursive Partitioning Using the RPART Routine. Technical Report 61. Section of Biostatistics, Mayo Clinic, Rochester.
- Venables, W.N., Ripley, B.D. (1994) *Modern Applied Statistics with S-PLUS*. Springer.
- Wright, G., Tan, B., Rosenwald, A., Hurt, E.H., Wiestner, A., Staudt, L.M. (2003) A gene expression-based method to diagnose clinically distinct subgroups of diffuse large B cell lymphoma. *Proc. Natl. Acad. Sci*, **100**(17), 9991-9996.