# DATA MINING IN FINANCIAL MARKETS

by

STEPHEN EVANS

Presented to the Faculty of the Graduate School of

The University of Texas at Arlington in Partial Fulfillment

of the Requirements

for the Degree of

MASTER OF SCIENCE IN INFORMATION SYSTEMS

THE UNIVERSITY OF TEXAS AT ARLINGTON

December 2011

## ACKNOWLEDGEMENTS

I would like to extend thanks to my supervising professor for the guidance and suggestions that he has given me in completing this study. I also thank him for being very instrumental in the design and implementation of this study. I would also like to thank the members of my committee along with my graduate advisor for their time and input in reviewing my work and offering criticism. I thank my mother for proofreading my paper. Finally, I thank my partner, Erica Robinson, for reviewing my paper and defense presentation.

November 22, 2011

ABSTRACT

DATA MINING IN FINANCIAL MARKETS

Stephen Evans, MS

The University of Texas at Arlington, 2011

Supervising Professor: Riyaz Sikora

Momentum in financial markets can cause securities prices to continue trending upward/downward based on the recent performance. This paper reviews a study that attempts to discover how much daily returns in the stock market can be explained by financial momentum. This study uses classification data mining to attempt to predict the direction of daily returns of randomly selected stocks from the Russell 1000 and Russell 2000 stock indexes. The study uses moving averages of historical daily stock prices as attributes, along with different data mining classifiers, to attempt to make these predictions. A secondary goal of this study is to determine how effective using Distributed Data Mining (DDM) can be in predicting the direction of daily stock returns. Hence, DDM classifiers are used in the testing.

This study discovers that the moving averages of daily returns do not help predict the direction of future daily stock returns any better than the percentages of returns from one trading day to the next. It also shows that the classifiers were no more than 60% accurate in predicting the directions of daily returns for any of the stocks used in this study. Hence, it appears that momentum cannot be used to explain very much of the movement in daily stock prices on a consistent basis.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

# CHAPTER 1

# INTRODUCTION

## 1.1 Overview of Data Mining

Data Mining is the process of sifting through data in order to find previously unrealized patterns (Sikora, 2011). This process is enabled by computer programs that are able to analyze various types of data using predetermined algorithms (Sikora, 2011). Data mining is/can be performed on data sets of many different types for the purpose of assisting individuals and/or organizations in making decisions.

There are 5 major types of data mining: Classification, Clustering, Association Rule Discovery, Regression, and Deviation Detection (Sikora, 2011). Classification is the most common predictive data mining type, and the one that is employed in this paper. It uses a subset of historical records from a given dataset with different attributes (this subset being referred to as the “training set”), to create a model for predicting the values of one of the attributes (referred to as the “class attribute”) using some combination of the remaining attributes. Based on how well this model performs in predicting the results of historical datasets that it has not yet seen, the model can be used to accurately predict future values of the class attribute (Sikora, 2011). These predictions, when accurate, can be used to help entities make more informed decisions regarding future events.

Distributed data mining (DDM) is a technique that can possibly be used to enhance the performance of existing classification models. While traditional classifiers use only one algorithm to attempt to explain all of the data in a given dataset, Distributed Learning Systems (DLS), use multiple algorithms to classify different subsets of the data in parallel (Sikora and Shaw, 1996). This allows for the best rules, from the best algorithm, to be matched with each of the examples, respectively (Sikora and Shaw, 1996). Potentially, using DDM can increase the performance of predicting examples in data mining.

### 1.2 Use of Data Mining in Financial Markets

One area in which data mining can be useful is in financial trading markets. In the U.S. Stock markets data mining can be used to help traders predict the outcomes of various securities. It can be used to find patterns in historical trading data that can be used by traders to garner better returns and make trading decisions with a much greater level of confidence. There has been a lot of research done already on data mining in financial markets. Many analyses have attempted to identify cross-correlations between different securities in various financial markets across the globe (Olson and Mossman, 2001 and Zhao et al., 2011 and Shapira et al., 2009). Some have attempted to identify correlations between financial securities and other possible predictors, such as social media or news feeds (Bourgon, 2010). Even others have attempted to predict the future results of financial securities using regression methods (Pollet and Wilson). These attempts have been met with varying degrees of success; while some attempts at accurately predicting future movements of financial securities have been moderately successful, many have not.

Of the attempts made so far to predict securities prices using historical data, few have attempted to predict these prices using classification data mining and derivations of a particular security's historical prices to predict future prices of that same security. This project attempts to do just that.

### 1.3 Financial Momentum

In financial markets the concept of momentum suggests that stock prices frequently continue to increase or decrease simply based on the trend in price movement from previous trading days (Crombez, 2001). Momentum in a stock's performance occurs when investors continue to buy a stock that is constantly increasing in value, or when investors continue to sell a stock that is constantly decreasing in value (Jegadeesh, 1993). This concept is thought to play a significant role in explaining the price movements in stock markets (Jegadeesh, 1993). This project attempts to test how much the concept of momentum can be used to explain the daily movement in stock prices.

CHAPTER 2

# INTRODUCTION

## 2.1 Goal of Project

The goal of this project is to find patterns between historical stock prices and future prices. More specifically, the goal of this project is to determine how much of the price fluctuation in the stock market can be explained by the concept of financial momentum. This research aims to accomplish this by attempting to predict the daily direction of movement for various stocks based on price data from previous trading days. A secondary goal of this study is to determine how effective using DDM can be in predicting the direction of daily stock returns in comparison to other classifier types.

## 2.2 Project Setup

### *2.2.1 Data Selection*

The data for this project was daily stock return data for stocks in the Russell 1000 and Russell 2000 stock indices. Ten stocks were randomly selected from each index so as to have representation of large and small market capitalization stocks in the tests. The stocks with an initial trading date after January 1$^{st}$ of 1996 were excluded from the pools of both indices before random sampling occurred in order to ensure that all of the stocks had trading data for the fifteen year period between January 1, 1996 and December 31, 2010. Each of the twenty stocks is identified by a single letter and a number (i.e. 'S1'). The letter (either 'S' or 'L') denotes whether the stock has a large market capitalization (L) or a small market capitalization (S). The ending number is a unique number within each market capitalization group.

### *2.2.2 Data Preparation*

The adjusted daily closing prices (the nominal closing prices for a stock modified to reflect any stock distributions or corporate actions that have taken place prior to the end of each day) for each of the stocks were used as the core attribute in trying to predict the performance of the stock (Investopedia, 2011). All of the other attributes used to help predict the daily direction of a particular stock's return were derived from the adjusted closing prices (Price). The additional attributes used for each stock are as follows: PriceDiff, SMA3, SMA3Diff, PriceDiffSMA3, SMA10, SMA10Diff, PriceDiffSMA10, SMA20, SMA20Diff, PriceDiffSMA20, SMA30, SMA30Diff, and PriceDiffSMA30. The attribute PriceDiff represents the percent difference of the adjusted close price of a given trading day and the adjusted closing price of the previous trading day. The SMA3 attribute represents the average of the closing prices of a particular trading day and its two previous trading days; the simple moving average of the close prices for 3 trading days. SMA3Diff is the percent difference of the current SMA3 and the previous SMA3, while the PriceDiffSMA3 represents the 3-day simple moving average of the PriceDiff. The remaining 10, 20, and 30-day attributes represent the same concepts as the corresponding 3-day attributes with the exception of the number of periods (days) used in calculating the simple moving averages. The final class attribute, NextDirection, indicates the direction of the next trading day's return for all trading days of a given stock.

### *2.2.3 Classifier Selection*

In this project several different classifiers are used to attempt to create a model for accurately predicting the direction of daily stock returns. These classifiers are grouped in the following categories: Rule-based, Decision Tree, Bayesian, Lazy, Multilayer Perceptron, and Meta.

Rule-based classifiers are classifiers in their most traditional sense (Sikora and Shaw, 1996). Rule-based classifiers are used to define rules that can be used to explain the results of a given training set (Sikora and Shaw, 1996). These classifiers begin with hypotheses of how the input attributes relate to the class attribute, and modify these hypotheses based on the data in the given training set (Sikora and Shaw, 1996). In this study OneR, Decision Table, and Ridor are rule-based classifiers that are used.

Decision tree classifiers use the input attributes from a dataset to create a tree structure used to predict the class attribute (Gehrke, 2003). A tree classifier iteratively partitions the input attributes into smaller parts (cells) that can be used to better predict the value of the class attribute (Hand, et al., 2001). Once the optimal partition sizes for each input attribute is determined, the cells can be used as nodes or leafs, to create a tree structure that can be used to predict the value of the class attribute (Hand, et al., 2001). J48 is a decision tree classifier that is used in this study.

Bayesian classifiers make use of conditional probability distributions to create a model to predict the class variable in a given dataset (Hand, 2001). This type of classifier initially calculates the probabilities that different values of given predictor attributes will correspond with any possible value of the class attribute (Madigan and Ridgeway, 2003). After these probabilities are gathered, the Bayesian classifiers use Bayes' theorem and the probabilities that were initially determined to calculate the conditional probabilities of each of the possible outcomes (Madigan and Ridgeway, 2003). Using the set of conditional probabilities, new observations in a dataset can be classified. This study uses the Naïve Bayes classifier.

Lazy classifiers are much simpler in concept than many of the other types of classifiers. These classifiers do not require a training phase (Sikora, 2011). They simply group examples in a dataset with its "nearest neighbors" or other examples that are most alike (Sikora, 2011). The variable KNN represents the maximum number of neighbors that should be contained in each group (Sikora, 2011). The weighted distance represents the maximum allowable distance between examples in the same group (Sikora, 2011). The weighted distance is more significant than the value of KNN in determining the size of a given group (Sikora, 2011). This study uses the lazy classifier: IBK.

A multilayer perceptron, one of the types of classifiers used in this study, is an artificial neural network based classifier (Hand, 2001). Multilayer Perceptron uses non-linear functions on the inputs (the predictor attributes) from a given dataset to predict the correct value of the class attribute (Si et al., 2003). Potentially, the final model will use a series of non-linear functions on the inputs in layers; using the initial inputs at the first layer and the results of the previous functions at subsequent layers (Hand, 2001). The end result is the prediction of the class attribute for a given example (Hand, 2001). During the training phase the type of functions that need to be run at each layer is determined; creating the optimal MLP model for the given dataset (Si et al., 2003).

The meta-classifiers used in this study are DDM, DDMwGA, Stacking, and Rotation Forest. These classifiers simply use a combination of already defined classifiers to predict the class attribute for a given dataset. Stacking uses a combination of user-defined classifiers while Rotation Forest uses a combination of decision tree classifiers to classify examples from a dataset. The DDM classifier in this study uses a combination of J48, Naïve Bayes, Multilayer Perceptron, IBK, and OneR to classify the examples. While DDMwGA is DDM with a genetic algorithm (GA), which facilitates the interaction of all the component classifiers in order to come up with the best possible classification for each example of a given dataset (Sikora and Shaw,1996).

## 2.3 Explanation of Tests

In the testing phase of this project various tests were run using the attributes described above in order to test the ability of these attributes to predict the direction of a stock's daily return. There are four core tests that were completed during this experiment. All of the tests included data from the 20 randomly chosen stocks (10 small cap and 10 large cap). Each of these tests were run more than once with minor adjustments made to them as thought to be necessary. However, only the four core tests will be explained in this section, as the revisions to these tests will be explained in the results chapter. The attributes for the stocks used in each test were broken down into separate attribute groups. Each attribute group was tested separately with the specifications of the given core test.

For the first test (Test 1) the examples for the trading days between January 1, 1996 and December 31, 2010 for each stock were taken as the subset used for testing. These attributes were rounded to 6 places after the decimal point. There were 6 separate attribute groups created for each stock. The groups are generally referred to as Basic, SMA3, SMA10, SMA20, SMA30, and All. These groups are not to be confused with the attributes themselves. The Basic attribute group includes the Price, and the PriceDiff attributes. The SMA3 includes the close price and all of the 3-day attributes. The SMA10, SMA20, and SMA30 attribute groups include the close price and all of their respective period attributes. The All attribute group includes all of the attributes used in this project. For all of the above 6 groups NextDirection is included with the values: 'Positive' (when the next trading day's return is zero or positive) and 'Negative' (when the next trading day's return is negative). Hence, the main focus in Test 1 is attempting to predict when the stock price return will be negative.

On the Test 1 datasets cross validation was used as the experiment type. Each set was tested with 10 folds. The classifiers used in Test 1 were J48, Naïve Bayes, Ridor, Multilayer Perceptron, Decision Table, OneR, and IBK (where KNN was set to two).

Test 2 was constructed and implemented exactly the same as Test 1, with the exception that the class attribute (NextDirection) had the following range of values: 'Positive' (when the next day's stock return is positive), 'Neutral' (when the next day's return is zero), and 'Negative (when the next day's stock return is negative). The focus of Test 2 is to test the accuracy of predicting the direction of daily returns when the daily stock return is negative, positive, or neutral.

Test 3 used the same test sets as Test 1. However, the Train/Test Percent Split experiment type was used in Test 3. The training percentage was set to 80. J48, Naïve Bayes, Multilayer Perceptron, OneR, and Rotation Forest were all used as classifiers for Test 3. In addition to these IBK was used where KNN was set to 100 and the distance weighting was set to 1/distance. The Stacking classifier was also used in this test, combining J48, Naïve Bayes, Multilayer Perceptron, OneR, and IBK (KNN =100, distance weighting = 1/distance), with J48 as the meta-classifier. Just as in Test 1 this test focuses on trying to predict negatively performing trading days.

Test 4 uses the same experiment type, datasets, and classifiers as Test 1 with only one exception. Test 4 focuses primarily on predicting positive trading days for a stock. Therefore, the range of the NextDirection values are: ‘Positive’ (where the next trading day’s return is positive) and ‘Negative’ (where the next trading day’s return is either negative or zero).

# CHAPTER 3

# INTRODUCTION

In this project there were 6 major variables that were expected to play a key role in determining the performance of predicting the directions of stock returns. These variables include the selection of attributes used in a test set, the classifiers used to test the data, the number of examples used in a test set, the particular stock, the market capitalization of the stock, and the method for calculating the NextDirection class attribute. Although it appears that some patterns may be present in the data, none of the results of these tests proved to be statistically or practically significant.

## 3.1 Analysis by Attribute Group

The attributes groups used in the tests run for this experiment proved to be insignificant in determining the performance of the prediction tests. However, there are some interesting observations that can be made from the results. The Basic and All attribute groups yielded the best results on average. As shown in Figure 3.1, the average of the percentages of NextDirection values that the classifiers accurately predicted was greatest when the attribute group was Basic or All. Since the Basic and All attribute groups have only the Basic predictor attributes in common, this may suggest that the Price and PriceDiff attributes are more useful in predicting the NextDirection attribute than any of the other predictor attributes. However, the differences between the percentages of correctly classified examples for the All and Basic attribute groups and those of the other attribute groups is rather small (within 1%).

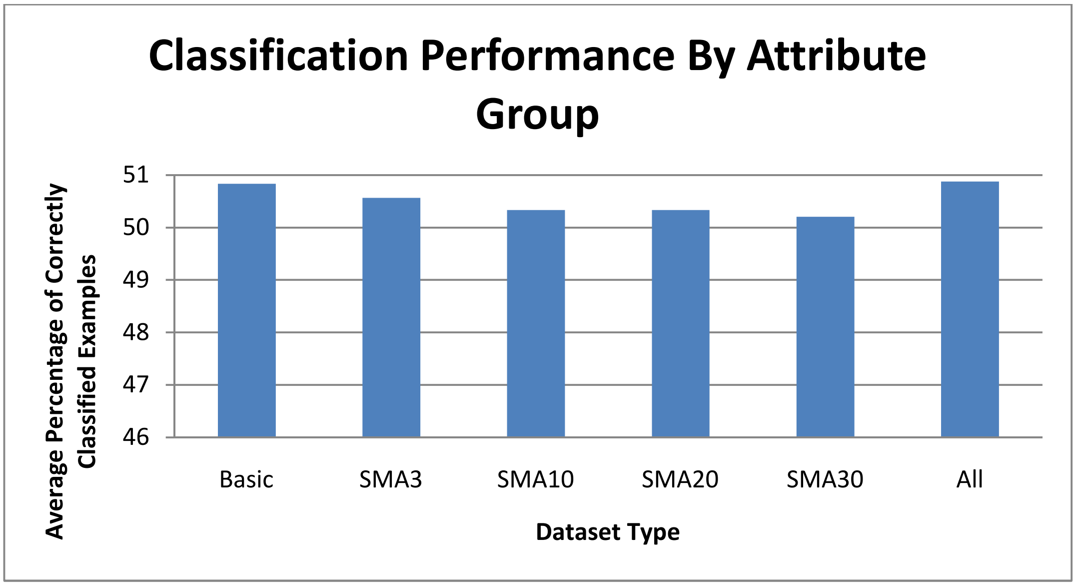


Figure 3.1 Percentages of Examples Correctly Classified by Attribute Group.

After its first run, Test 1 was revised (as Test 1 Revised) to try and determine the effect that removing the most basic attributes, Price and PriceDiff, from all of the datasets would have on the performance of predicting the direction of the next period's Price. This involved excluding the Basic attribute group altogether, removing the Price attribute from the SMA attribute groups, and removing the Price and PriceDiff attributes from the All attribute group. The results of Test 1 Revised (Figure 3.2) were almost exactly the same as the Test 1 results, for the SMA attribute groups. There is no comparison for the Basic dataset between these two tests, since the removal of the Price and PriceDiff attributes in Test 1 Revised eliminates the entire Basic attribute group. However, the results for the All attribute group were noticeably larger in Test 1 Revised than in Test 1. This seems to take away from the idea that the predictor attributes in the Basic attribute set are more instrumental in helping to predict the class attribute than the other attributes.

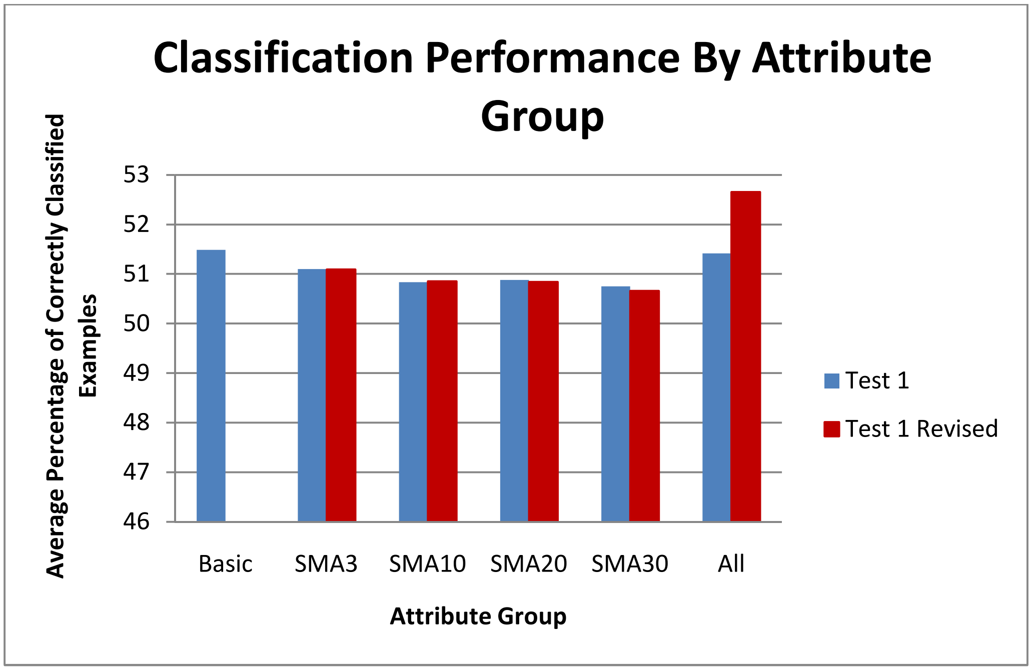


Figure 3.2 Percentage of Examples Correctly Classified in Test1 and Test 1 Revised by Attribute Group.

## 3.2 Analysis by Classifier

Among the 4 core tests, J48 and Decision Table appear to be the two best performing classifiers that were consistently used. In Figure 3.3 J48 and Decision Table are shown as the classifiers with the best two percentages of correctly classified examples. This same trend held true for each of the 4 tests individually with the exception of Test 3 (shown in Figure 3.4), where Decision Table was not used, and Rotation Forest and J48 were the two best performing classifiers. Also, J48 and Decision Table were the only two classifiers to perform noticeably better than ZeroR (which should produce the same results as random guessing) in Figure 3.3. Rotation Forest, J48, and Naïve Bayes were the only three to perform better than ZeroR, on average, in Test 3.

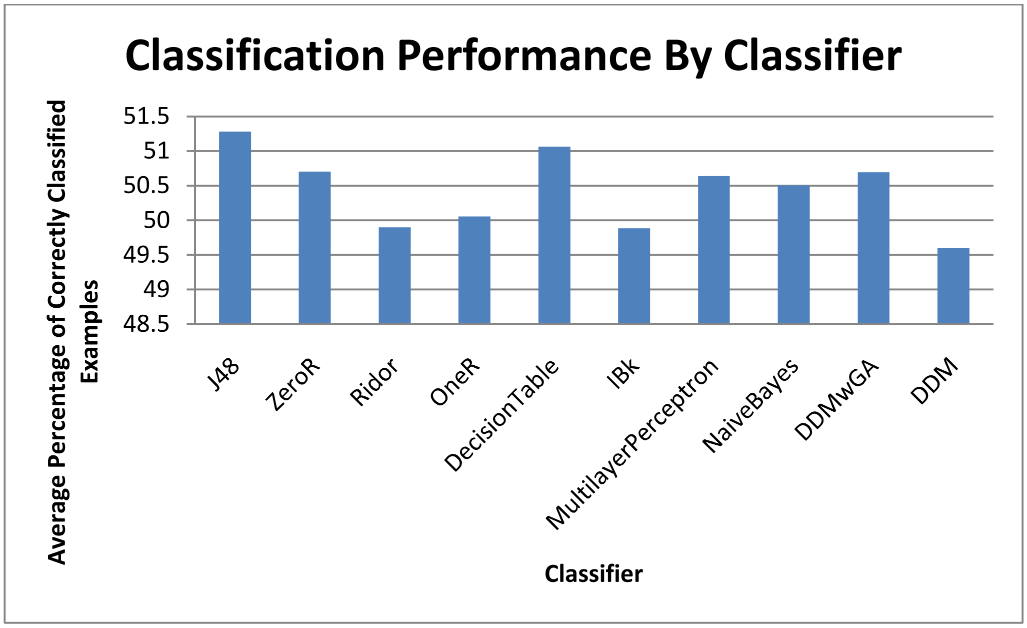


Figure 3.3 Percentages of Examples Correctly Classified by Classifier.

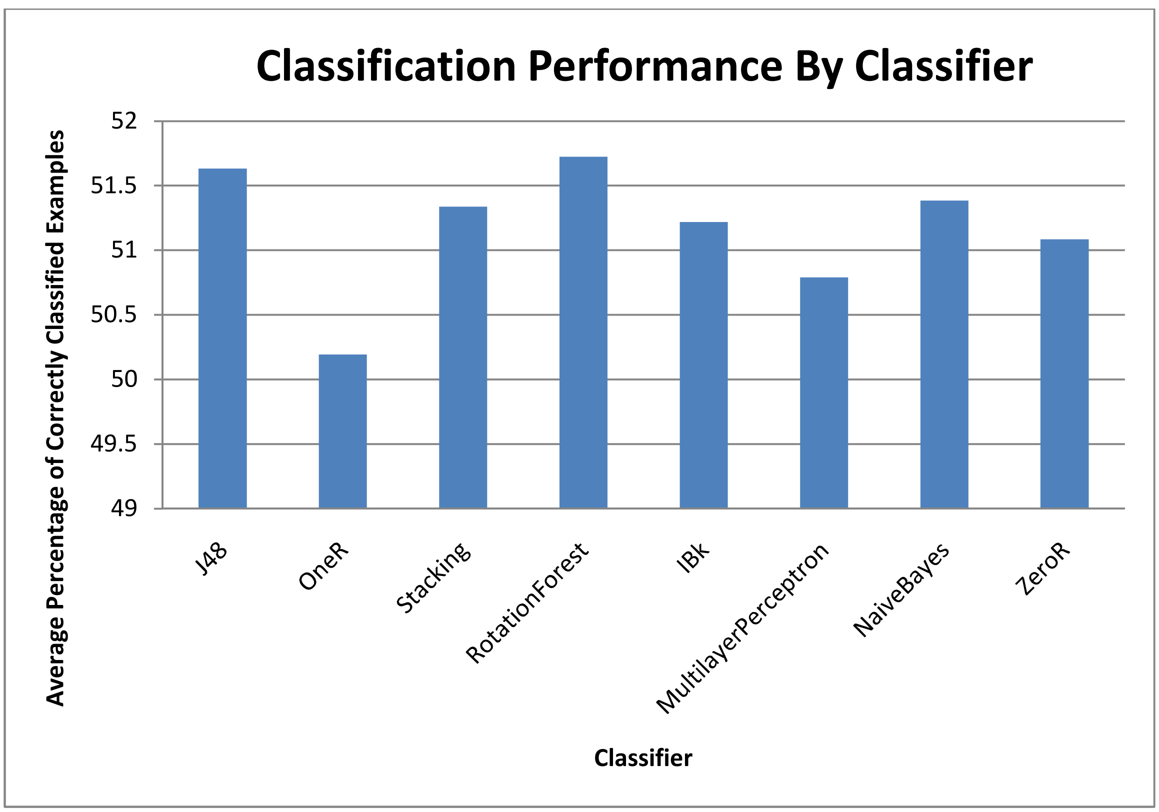


Figure 3.4 Percentages of Examples Correctly Classified in Test3 by Classifier.

The results of these tests might also seem to indicate that tree learners can perform among the better classifiers in predicting the direction of stock returns since J48 is a decision tree and Rotation Forest is a meta-classifier consisting of tree learners. Also, while DDMwGA seemed to be a better performer in predicting the class attribute than DDM, neither one of these attributes appear to perform better than ZeroR, Decision Table, or J48 on average.

On the other hand, when looking at the results of these tests by stock and classifier (Appendix A) it is easy to see that the generalizations above don't necessarily hold true for each stock. Although J48 and Decision Table rank among the top two classifiers for a given stock more frequently than any of the other classifiers in Appendix A, there are several stocks for which they do not. For a majority of these stocks the differences between the results of the classifiers are very small and/or the classifier results are less than or not much larger than random guessing. However, the results for a few stocks appear to merit some mention. For example Naïve Bayes seems to be better at predicting the direction of daily returns for L2 and L5. Also, for S9 DDMwGA, although not noticeably better than Decision Table or J48, appears to be among the better classifiers for predicting the class attribute.

## 3.3 Analysis by Size of Subset

Test 1 and Test 3 were modified to include all of the stock price data for the 20 randomly chosen stocks. For each stock in these modified tests (Test 1 Fulltime and Test 3 Fulltime) the trading data from the stock's first trading date to the beginning of this project was used in each dataset. All of the other variables from the original tests were held constant. From the results of these tests it appears that (on average) the number of trading days used in testing is likely not relevant to the ability of the classifiers to predict the direction of the stock returns. Although the average percentage of correctly classified samples increased from the original test to the fulltime test for Test 1 and Test 3, the differences in the averages were less than 0.4% in both cases.

However, this seems to affect individual stocks differently. In Appendix B, which shows the results of Test 1, Test 1 Fulltime, Test 3, and Test 3 Fulltime, by stock (including the first trading date of the stock), there is no apparent trend in the performance as the initial trading date progresses. Yet, there appears to be a noticeable increase for S1 and L9 from the core tests to the fulltime tests, suggesting that the size of the dataset may have a greater effect on the performance of predicting the direction of daily returns for these stocks than it does for others.

## 3.4 Analysis by Market Capitalization

Market Capitalization of the stocks did not seem to play much of a role in how well the classifiers performed in predicting the daily direction of a stock's movement. On average, small cap stocks performed better than large cap stocks, but with a margin of half of one percent. The daily return directions for small cap stocks were more accurately predicted (on average) for three of the 4 core tests than those for the large cap stocks, but the margins were not very large. Even the results as shown per classifier (Figure 3.5) show little difference in classification performance between small and large cap stocks. Although, Figure 3.5 seems to indicate that Decision Table and J48 were able to gain much of their advantage over the other classifiers with their performances in predicting the class attribute for the small cap stocks.

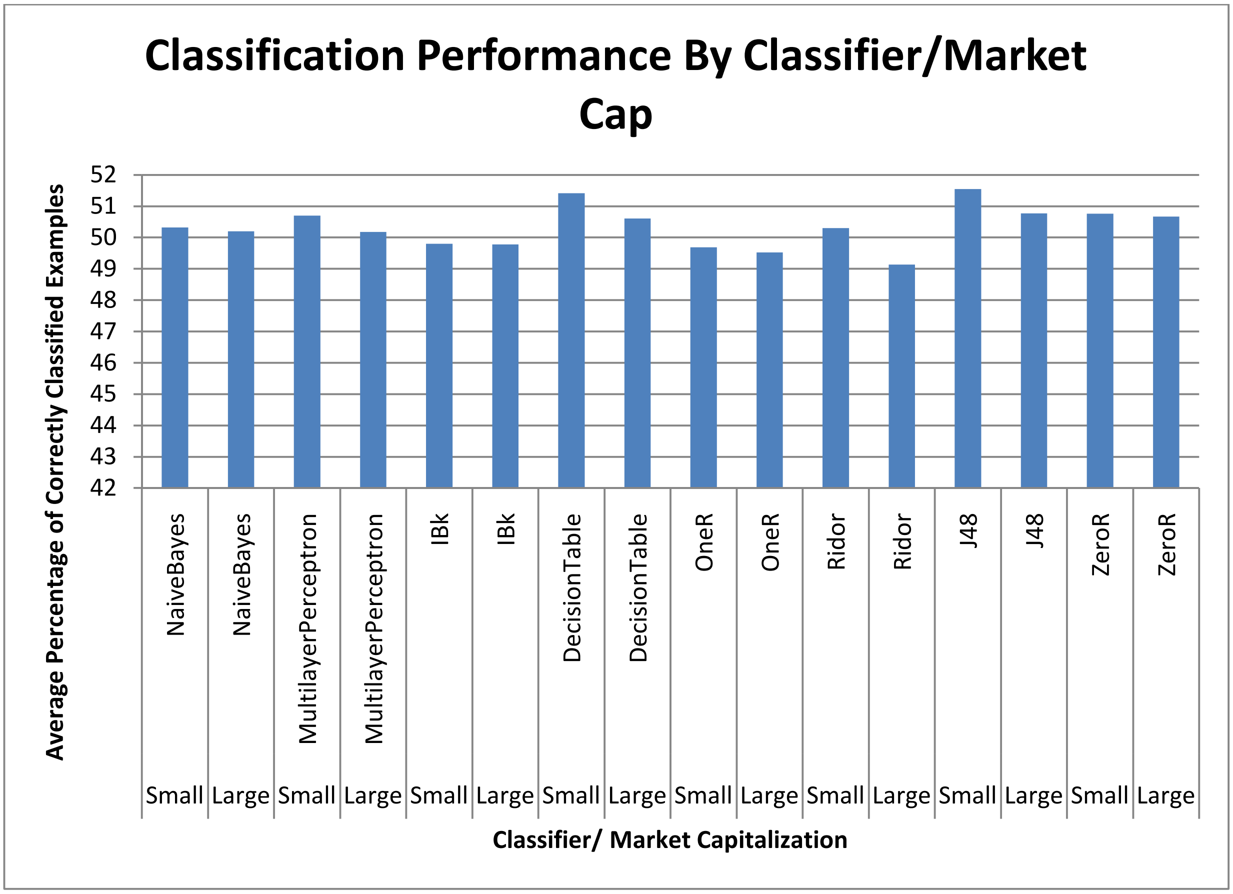


Figure 3.5 Classification Performance Per Classifier Per Market Capitalization.

## 3.5 Analysis by Stock

The performance of classifying the class attribute, on average, did not vary much by stock, although classifications of the examples for a few stocks appear to be more accurate than classifications for the other stocks. In Figure 3.6 the examples for S1, and S10 appear to (on average) be better classified than those for the other stocks. These two stocks are noticeably higher on the graph than the other stocks, while the remaining stocks in this figure are all within about 1 percent of each other.

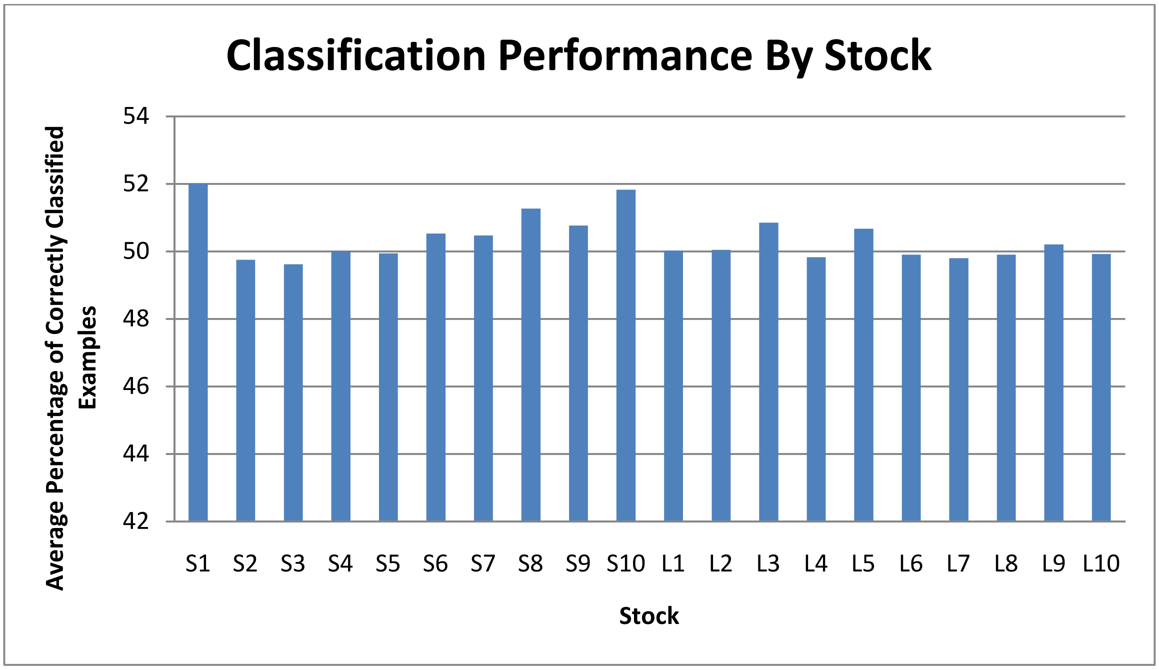


Figure 3.6 Classification Performance By Stock.

## 3.6 Analysis by Class Attribute Calculation

The method of calculation used for the class attribute seems to play a part in how well the classifiers perform in predicting daily stock return directions. On average the percentage of correctly classified examples decreases from Test 1 (51.08%, where the NextDirection attribute is focused on negative directions) to Test 2 (47.10%, where NextDirection includes ‘Negative’, ‘Positive’, and ‘Neutral’ values), and increases from Test 1 to Test 4 (52.10%, where NextDirection is focused on positive returns). Hence, it appears that the classifiers may perform better at classifying positive returns than when attempting to predict negative returns or all three directions simultaneously.

Figure 3.7, which shows the performance values for predicting the class attribute from Test 1, Test 2, and Test 4 by stock, seems to partially support the idea that the class attribute in Test 1 can be more accurately classified than the class attributes in Test 2 and Test 4. The graph in this figure depicts Test 2 as noticeably being the worst performing test for all of the stocks, which may not be much of a surprise since Test 2 has a 3-class class attribute while

Test 1 and Test 4 have 2-class class attributes. However, the graph also shows that the Test 4 performance is greater than the Test 1 performance for the majority of the stocks. Yet, Test 4 results are more than 1 percent larger than the corresponding Test 1 results for only a subset of those stocks.

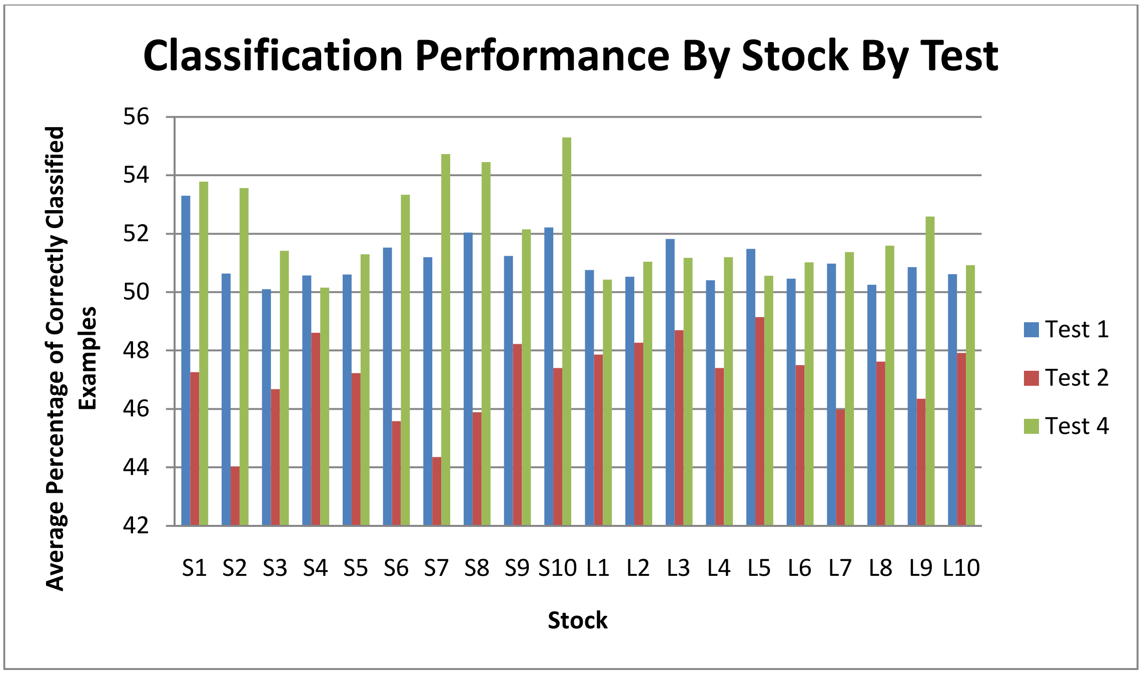


Figure 3.7 Classification Performance Per Stock Per Class Attribute Calculation.

# CHAPTER 4

# INTRODUCTION

## 4.1 Summary of Findings

The results of this study turned out not to be significant, either statistically or practically. However, there are a few interesting observations worth noting from the results. Decision Table, J48, and Rotation Forest appear to be the best of all the classifiers used in predicting the direction of daily stock returns. Although this could not be confirmed statistically, it may be able to be confirmed with further testing. The size of the dataset used may be significant in predicting the direction of future returns, but if so, this will likely depend on the stock. The directions of daily returns of stocks with small market capitalizations can possibly be more accurately predicted than those of stocks with large market capitalizations, though this will likely differ from one stock to the next. The method for setting the NextDirection attribute might also be significant in classification performance. It appears that the positive daily returns could possibly be better predicted than negative ones. Finally, contrary to initial expectations the attribute group does not seem to have a clear effect on the classification performance for daily stock return directions.

## 4.2 Measure of Success in Achieving Goals

Due to the fact that the results of this study were not statistically significant, success in achieving the goals of this project cannot be fully claimed. The performance of predicting the direction of daily stock returns was far from adequate for practical use in all of the tests in this project. None of the prediction accuracies were much larger (if larger at all) than the expected accuracy of random guessing. Hence, the primary goal of this study has been achieved with the conclusion that none of the performance of daily stock returns used in this study can really be attributed to financial momentum. However, the secondary goal of determining how well Distributed Data Mining classifiers would perform against other classifiers has not been realized in this project. The results appear to indicate that DDM and DDMwGA classifiers perform substandard to J48, Decision Table, and possibly Rotation Forest, and that each of the classifiers may perform differently in comparison to the others based on the stock for the returns being predicted. Yet, none of these results are significant and therefore it cannot be determined how well DDM classifiers perform in comparison with the other classifiers used in this study.

### 4.3 Ideas for Future Testing

To obtain better results and fully achieve the goals of this study in the future, there are some changes that can be made to the study. The performance of predicting the daily returns of stocks may benefit from a smaller or more targeted dataset size. In other words, returns during certain economic recession periods may be better predicted in this study than an uncategorized multiyear range of returns. Removing outliers from the datasets may also result in better performance in predicting daily stock returns. Using weighted averages among the predictor attributes may also increase the performance of classification. Also, since the concept of financial momentum is more commonly associated with longer term trading, increasing the number of days used in the moving averages or increasing the size of the return period may produce better results.

# APPENDIX A

## CLASSIFICATION PERFORMANCE PER STOCK PER CLASSIFIER

| Stock | Classifier | Test Average |
|---|---|---|
| S1 | DDM | 49.40 |
| S1 | DDMwGA | 52.04 |
| S1 | NaiveBayes | 50.08 |
| S1 | MultilayerPerceptron | 50.78 |
| S1 | IBk | 48.74 |
| S1 | DecisionTable | 52.38 |
| S1 | OneR | 48.61 |
| S1 | Ridor | 50.77 |
| S1 | ZeroR | 50.00 |
| S1 | J48 | 52.27 |
| L1 | DDM | 49.56 |
| L1 | DDMwGA | 49.66 |
| L1 | NaiveBayes | 48.23 |
| L1 | MultilayerPerceptron | 50.15 |
| L1 | IBk | 49.76 |
| L1 | DecisionTable | 50.24 |
| L1 | OneR | 49.01 |
| L1 | Ridor | 46.54 |
| L1 | ZeroR | 50.27 |
| L1 | J48 | 50.09 |
| L2 | DDM | 49.52 |
| L2 | DDMwGA | 50.72 |
| L2 | NaiveBayes | 51.30 |
| L2 | MultilayerPerceptron | 50.22 |
| L2 | IBk | 48.84 |
| L2 | DecisionTable | 50.24 |
| L2 | OneR | 49.08 |
| L2 | Ridor | 47.77 |
| L2 | ZeroR | 50.24 |
| L2 | J48 | 50.14 |

| Stock | Classifier | Test Average |
|---|---|---|
| S2 | DDM | 48.41 |
| S2 | DDMwGA | 49.69 |
| S2 | NaiveBayes | 48.54 |
| S2 | MultilayerPerceptron | 48.80 |
| S2 | IBk | 46.84 |
| S2 | DecisionTable | 50.11 |
| S2 | OneR | 47.95 |
| S2 | Ridor | 49.48 |
| S2 | ZeroR | 50.00 |
| S2 | J48 | 49.85 |
| S3 | DDM | 49.07 |
| S3 | DDMwGA | 50.45 |
| S3 | NaiveBayes | 48.63 |
| S3 | MultilayerPerceptron | 49.60 |
| S3 | IBk | 48.88 |
| S3 | DecisionTable | 50.14 |
| S3 | OneR | 48.82 |
| S3 | Ridor | 47.34 |
| S3 | ZeroR | 50.24 |
| S3 | J48 | 49.93 |
| L3 | DDM | 49.33 |
| L3 | DDMwGA | 50.60 |
| L3 | NaiveBayes | 49.49 |
| L3 | MultilayerPerceptron | 50.62 |
| L3 | IBk | 49.45 |
| L3 | DecisionTable | 50.11 |
| L3 | OneR | 49.45 |
| L3 | Ridor | 49.44 |
| L3 | ZeroR | 50.34 |
| L3 | J48 | 51.01 |

| Stock | Classifier | Test Average |
|---|---|---|
| S4 | DDM | 49.14 |
| S4 | DDMwGA | 49.23 |
| S4 | NaiveBayes | 48.71 |
| S4 | MultilayerPerceptron | 49.41 |
| S4 | IBk | 49.51 |
| S4 | DecisionTable | 49.89 |
| S4 | OneR | 48.94 |
| S4 | Ridor | 49.32 |
| S4 | ZeroR | 50.00 |
| S4 | J48 | 49.92 |
| S5 | DDM | 49.01 |
| S5 | DDMwGA | 49.33 |
| S5 | NaiveBayes | 49.01 |
| S5 | MultilayerPerceptron | 49.39 |
| S5 | IBk | 48.41 |
| S5 | DecisionTable | 49.55 |
| S5 | OneR | 49.04 |
| S5 | Ridor | 49.65 |
| S5 | ZeroR | 50.00 |
| S5 | J48 | 49.78 |
| L4 | DDM | 49.08 |
| L4 | DDMwGA | 49.69 |
| L4 | NaiveBayes | 49.35 |
| L4 | MultilayerPerceptron | 49.31 |
| L4 | IBk | 49.48 |
| L4 | DecisionTable | 50.00 |
| L4 | OneR | 48.80 |
| L4 | Ridor | 48.42 |
| L4 | ZeroR | 50.00 |
| L4 | J48 | 49.72 |

| Stock | Classifier | Test Average |
|---|---|---|
| L5 | DDM | 49.47 |
| L5 | DDMwGA | 50.62 |
| L5 | NaiveBayes | 51.06 |
| L5 | MultilayerPerceptron | 50.31 |
| L5 | IBk | 49.10 |
| L5 | DecisionTable | 50.92 |
| L5 | OneR | 49.25 |
| L5 | Ridor | 47.48 |
| L5 | ZeroR | 50.00 |
| L5 | J48 | 50.84 |
| S6 | DDM | 49.31 |
| S6 | DDMwGA | 50.43 |
| S6 | NaiveBayes | 47.70 |
| S6 | MultilayerPerceptron | 50.63 |
| S6 | IBk | 48.38 |
| S6 | DecisionTable | 51.22 |
| S6 | OneR | 48.37 |
| S6 | Ridor | 49.03 |
| S6 | ZeroR | 50.00 |
| S6 | J48 | 50.88 |
| S7 | DDM | 49.15 |
| S7 | DDMwGA | 50.85 |
| S7 | NaiveBayes | 48.10 |
| S7 | MultilayerPerceptron | 51.53 |
| S7 | IBk | 47.32 |
| S7 | DecisionTable | 51.50 |
| S7 | OneR | 48.66 |
| S7 | Ridor | 48.31 |
| S7 | ZeroR | 51.38 |
| S7 | J48 | 51.38 |

| Stock | Classifier | Test Average |
|---|---|---|
| L6 | DDM | 49.28 |
| L6 | DDMwGA | 50.07 |
| L6 | NaiveBayes | 49.86 |
| L6 | MultilayerPerceptron | 49.64 |
| L6 | IBk | 49.26 |
| L6 | DecisionTable | 50.00 |
| L6 | OneR | 49.28 |
| L6 | Ridor | 46.58 |
| L6 | ZeroR | 50.00 |
| L6 | J48 | 50.21 |
| L7 | DDM | 48.32 |
| L7 | DDMwGA | 49.34 |
| L7 | NaiveBayes | 47.88 |
| L7 | MultilayerPerceptron | 49.03 |
| L7 | IBk | 46.65 |
| L7 | DecisionTable | 50.02 |
| L7 | OneR | 48.15 |
| L7 | Ridor | 49.17 |
| L7 | ZeroR | 50.00 |
| L7 | J48 | 49.82 |
| S8 | DDM | 49.56 |
| S8 | DDMwGA | 50.85 |
| S8 | NaiveBayes | 49.05 |
| S8 | MultilayerPerceptron | 50.49 |
| S8 | IBk | 49.76 |
| S8 | DecisionTable | 51.65 |
| S8 | OneR | 49.22 |
| S8 | Ridor | 50.00 |
| S8 | ZeroR | 50.19 |
| S8 | J48 | 51.05 |

| Stock | Classifier | Test Average |
|---|---|---|
| S9 | DDM | 49.36 |
| S9 | DDMwGA | 51.07 |
| S9 | NaiveBayes | 50.88 |
| S9 | MultilayerPerceptron | 50.02 |
| S9 | IBk | 48.70 |
| S9 | DecisionTable | 51.06 |
| S9 | OneR | 49.67 |
| S9 | Ridor | 49.94 |
| S9 | ZeroR | 50.00 |
| S9 | J48 | 51.05 |
| L8 | DDM | 49.03 |
| L8 | DDMwGA | 49.62 |
| L8 | NaiveBayes | 49.24 |
| L8 | MultilayerPerceptron | 49.90 |
| L8 | IBk | 50.75 |
| L8 | DecisionTable | 49.77 |
| L8 | OneR | 48.62 |
| L8 | Ridor | 48.60 |
| L8 | ZeroR | 50.52 |
| L8 | J48 | 50.36 |
| L9 | DDM | 48.80 |
| L9 | DDMwGA | 50.63 |
| L9 | NaiveBayes | 48.01 |
| L9 | MultilayerPerceptron | 50.93 |
| L9 | IBk | 46.99 |
| L9 | DecisionTable | 50.99 |
| L9 | OneR | 49.01 |
| L9 | Ridor | 49.56 |
| L9 | ZeroR | 50.00 |
| L9 | J48 | 50.79 |

| Stock | Classifier | Test Average |
|---|---|---|
| S10 | DDM | 50.50 |
| S10 | DDMwGA | 52.55 |
| S10 | NaiveBayes | 49.50 |
| S10 | MultilayerPerceptron | 52.51 |
| S10 | IBk | 47.27 |
| S10 | DecisionTable | 53.48 |
| S10 | OneR | 49.80 |
| S10 | Ridor | 53.04 |
| S10 | ZeroR | 53.10 |
| S10 | J48 | 53.85 |
| L10 | DDM | 49.44 |
| L10 | DDMwGA | 50.15 |
| L10 | NaiveBayes | 48.42 |
| L10 | MultilayerPerceptron | 48.93 |
| L10 | IBk | 49.94 |
| L10 | DecisionTable | 50.64 |
| L10 | OneR | 49.19 |
| L10 | Ridor | 48.22 |
| L10 | ZeroR | 50.53 |
| L10 | J48 | 50.59 |

# APPENDIX B

## CLASSIFICATION PERFORMANCE PER STOCK/INITIAL TRADING DATE

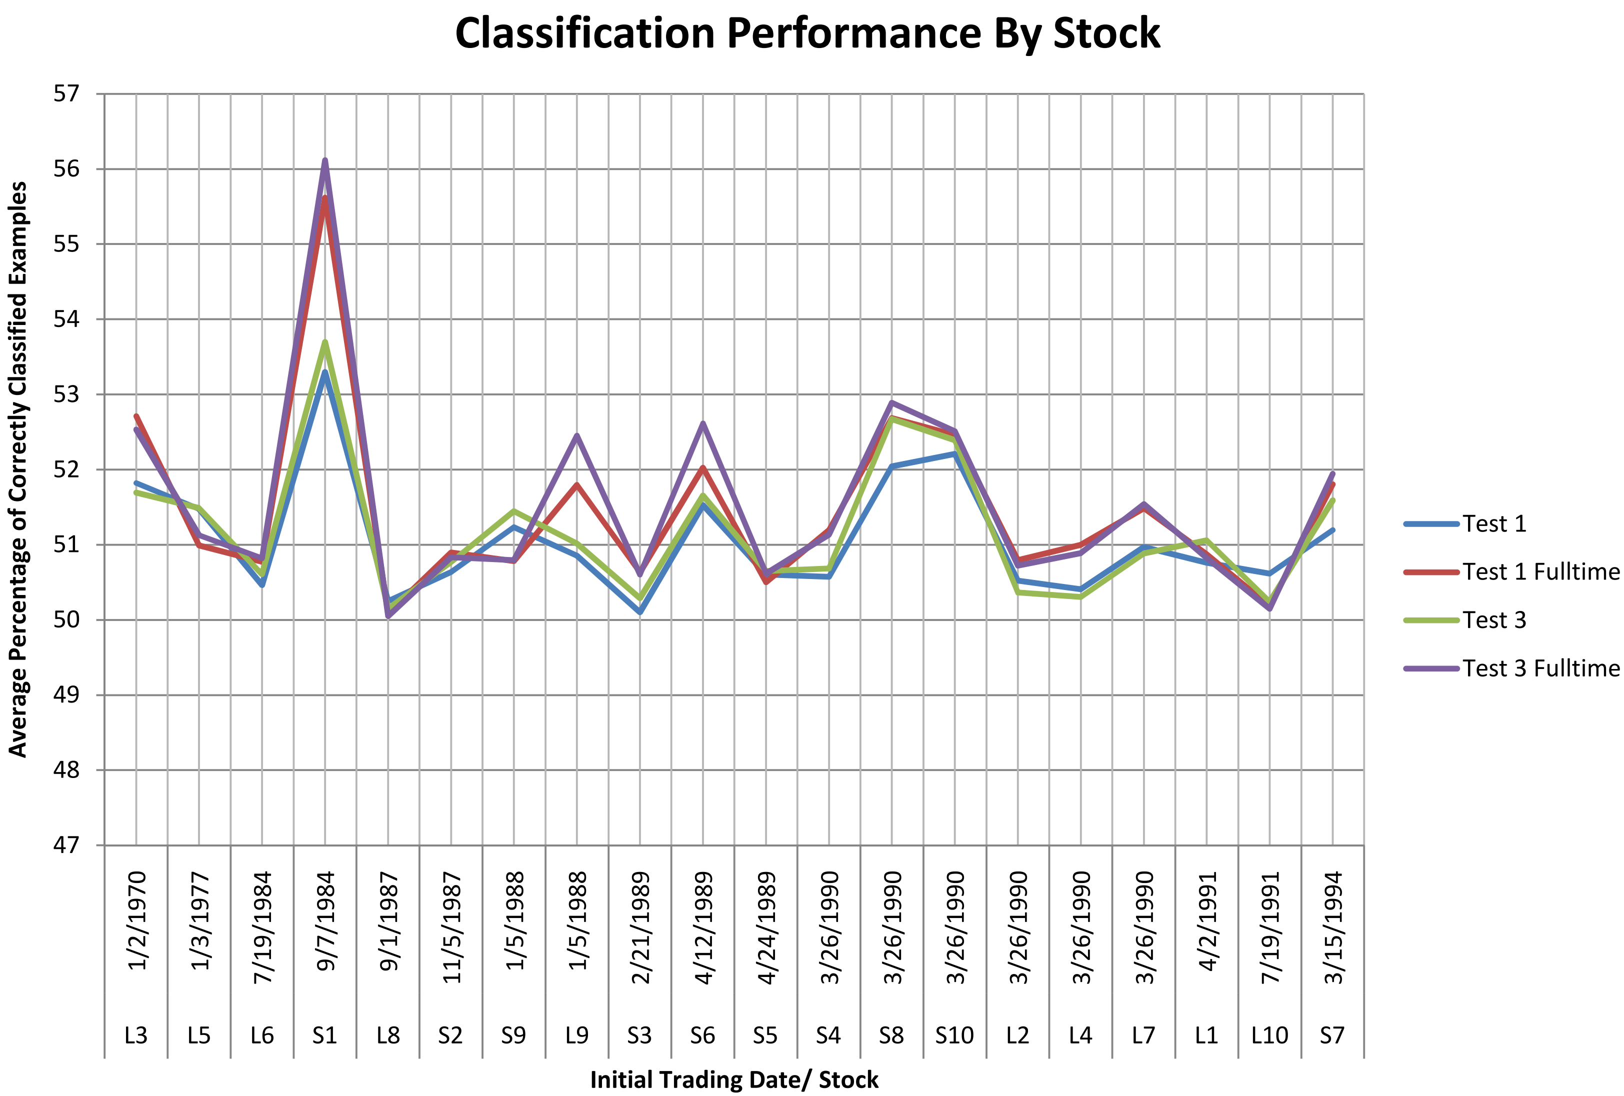
Classification Performance By Stock
Average Percentage of Correctly Classified Examples
57
56
55
54
53
52
51
50
49
48
47
1/2/1970 L3
1/3/1977 L5
7/19/1984 L6
9/7/1984 S1
9/1/1987 L8
11/5/1987 S2
1/5/1988 S9
1/5/1988 L9
2/21/1989 S3
4/12/1989 S6
4/24/1989 S5
3/26/1990 S4
3/26/1990 S8
3/26/1990 S10
3/26/1990 L2
3/26/1990 L4
3/26/1990 L7
4/2/1991 L1
7/19/1991 L10
3/15/1994 S7
Initial Trading Date/ Stock
Test 1
Test 1 Fulltime
Test 3
Test 3 Fulltime

REFERENCES

"Adjusted Closing Price." *Investopedia.com*. Investopedia ULC, 2011. Web. 2 Dec. 2011.

Bourgon, Lyndsie. "A Mood Ring for the Stock Market." *Canadian Business* 83.19 (2010): 26.

Crombez, J. "Momentum, Rational Agents and Efficient Markets." *The Journal of Psychology and Financial Markets*, 2, 2001.

Gehrke, Johannes. "Decision Trees." *The Handbook of Data Mining.* Ed. Nong Ye. Mahwah: Lawrence Erlbaum Associates, Inc. Publishers, 2003.3-24.

Hand, David, Heikki Mannila, and Padhraic Smyth. *Principles of Data Mining*. Cambridge, Massachusetts: The MIT Press, 2001. Print.

Jegadeesh, N., and S. Titman. "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency." *Journal of Finance*, 48, 1993.

Madigan, David, and Greg Ridgeway. "Bayesian Data Analysis." *The Handbook of Data Mining.* Ed. Nong Ye. Mahwah: Lawrence Erlbaum Associates, Inc. Publishers, 2003. 104-131.

Olson, D., and C. Mossman. "Cross-correlations and Predictability of Stock Returns." *Journal of Forecasting* 20.2 (2001): 145-160.

Pollet, Joshua M., and Mungo Wilson. "Average correlation and stock market returns." *Journal of Financial Economics* 96.3 (2010): 364-380.

Shapira, Y., D. Y. Kenett, and E. Ben-Jacob. "The Index cohesive effect on stock market correlations." *European Physical Journal B -- Condensed Matter* 72.4 (2009): 657-669.

Si, Jennie, Benjamin J. Nelson, and George C.Runger. "Artificial Neural Network Models for Data Mining." *The Handbook of Data Mining.* Ed. Nong Ye. Mahwah: Lawrence Erlbaum Associates, Inc. Publishers, 2003. 41-66.

Sikora, Riyaz. Class Lecture. Business Data Mining. University Texas at Arlington, Arlington, TX. 27 Jan 2011.

Sikora, Riyaz and Michael Shaw. "A Computational Study of Distributed Rule Learning." *Information Systems Research* 7.2 (1996): 189-197.

Zhao, Xiaojun, Pengjian Shang, and Qiuyue Jin. "Multifractal Detrended Cross-Correlation Analyis of Chiness Stock Markets Based on Time Delay." *Fractals* 19.3 (2011): 329-338.

# BIOGRAPHICAL INFORMATION

Stephen Evans enrolled in Xavier University as an undergraduate physics major in 2002. In the summer of 2004 Stephen worked on a materials science project at the University of Alabama at Birmingham, which explored applying diamond coatings to Nitinol (nickel-titanium) stents in an attempt to increase their lifespan. In 2005 Stephen transferred to the University of Texas at Arlington where his interest shifted from physics to information systems. In the spring of 2009 Stephen graduated cum laude with a Bachelor of Science degree in Information Systems from the University of Texas at Arlington. Stephen will complete his Masters of Science in Information Systems from the University of Texas at Arlington in December of 2011. After he graduates, Stephen plans to continue to work as an IT consultant. One of Stephen's long term plans is to work on projects that combine his interest in finance and technology.