A multi-component nonlinear prediction system for the S & P 500 Index

Tim Chenoweth [a,b,c,*], Zoran Obradović [a,*]

[a] *School of Electrical Engineering and Computer Science, Washington State University, Pullman WA 99164-2752, USA*
[b] *Department of Management and Systems, Washington State University, Pullman WA 99164-2752, USA*
[c] *Department of Economics, Washington State University, Pullman WA 99164-2752, USA*

Abstract

The proposed stock market prediction system is comprised of two preprocessing components, two specialized neural networks, and a decision rule base. First, the preprocessing components determine the most relevant features for stock market prediction, remove the noise, and separate the remaining patterns into two disjoint sets. Next, the two neural networks predict the market's rate of return, with one network trained to recognize positive and the other negative returns. Finally, the decision rule base takes both return predictions and determines a buy/sell recommendation. Daily and monthly experiments are conducted and performance measured by computing the annual rate of return and the return per trade. Comparison of the results achieved by the dual neural network system to that of the single neural network shows that the dual neural network system gives much larger returns with fewer trades. In addition, dual neural network experiments with the appropriately selected filtering and decision thresholds managed to achieve an almost twice larger annual rate of return when compared to that of the buy and hold strategy over a seventy month period. However, no claims are made that the proposed system is better than the buy and hold strategy when considering transaction costs.

Keywords: Stock market prediction; Hierarchical systems; Hybrid systems; Neural networks

1. Introduction

Financial markets in general and the stock market in particular are extremely efficient, meaning that at any given point in time the market does a very good job

---

* Corresponding author. Email: zoran@eecs.wsu.edu

reflecting the actual value of the underlying stocks. In fact, the efficient market hypotheses states that any new information which affects this value is accounted for by the market before the general public can make trades based on it [13]. This hypothesis seems reasonable *provided* that the relationship between the new information and the value of the underlying stock is fully understood. Most quantitative methods attempting to capture such relationships are based on simple stochastic linear time series models [2,14,16]. This research, similar to [12], hypothesizes that there might be nonlinear relationships between market information and the value of stocks that so far have not been identified and therefore are not reflected in stock prices. In other words, there might still be inefficiencies in the market. It is important to note that if this is true and nonlinear relationships do exist, once their existence (and the means for identifying them) becomes public knowledge, traders will have a better understanding of the relationship between information and stock values, which means both a more efficient market and the elimination of these nonlinear relationships. The long term results then will not be higher returns, but a more efficient market.

If these nonlinear relationships exist, it may be possible to capture them using a nonparametric machine learning approach of multilayer artificial neural networks (NN). Such NNs are powerful computational systems that theoretically can approximate any nonlinear continuous function on a compact domain to any desired degree of accuracy [9]. In addition, a NN can account for fundamental changes in the underlying function through incremental retraining. However, the stock market is a highly complex system which takes a very large quantity of information (fundamentals, news, etc.) and produces a price movement. What a NN is trying to do is model this system using only a very small portion of the total information. The other information not accounted for now plays the role of noise, resulting in a problem domain that is extremely noisy. In addition, due to the continuous evolution of the market, historical data may represent patterns of behavior that no longer hold [1].

One approach to modeling the stock market is to preprocess the data using formal feature selection and noise removal techniques and then train a NN on the preprocessed data. However, results from this study indicate that the combination of a noisy environment and non-stationarity makes it very difficult for a single NN to learn a function which generalizes when applied to new data. While the data preprocessing did improve the results of the single NN system, the system still did not achieve a satisfactory annual rate of return when compared to that achieved using the simple buy and hold strategy. Therefore, this paper proposes a hybrid multi-component nonlinear system for S&P 500 stock market predictions that utilizes two NNs, one trained on patterns corresponding to up movements in the market and the other trained on patterns corresponding to down movements.

The system goals are to earn a larger annual return than the buy and hold strategy and to keep the number of trades low to reduce transaction costs. The system details are explained in Section 2 followed by results and analysis in Section 3, and conclusions in Section 4. An extended abstract of this paper using only daily

data appears in [6]. The reader is also referred to companion papers [5,7] dealing with the specific issue of feature selection.

2. Methodology

The proposed system consists of a statistical feature selection component for identification of the most relevant data, a data filtering component for removing noise and splitting the remaining data, two specialized NNs for extraction of nonlinear relationships from the selected data, and high level decision rules for determining buy/sell recommendations (see Fig. 1).

2.1. Feature selection

The objective of the feature selection component is to identify a small subset of the most relevant features from a larger pool for designing the system in a manner
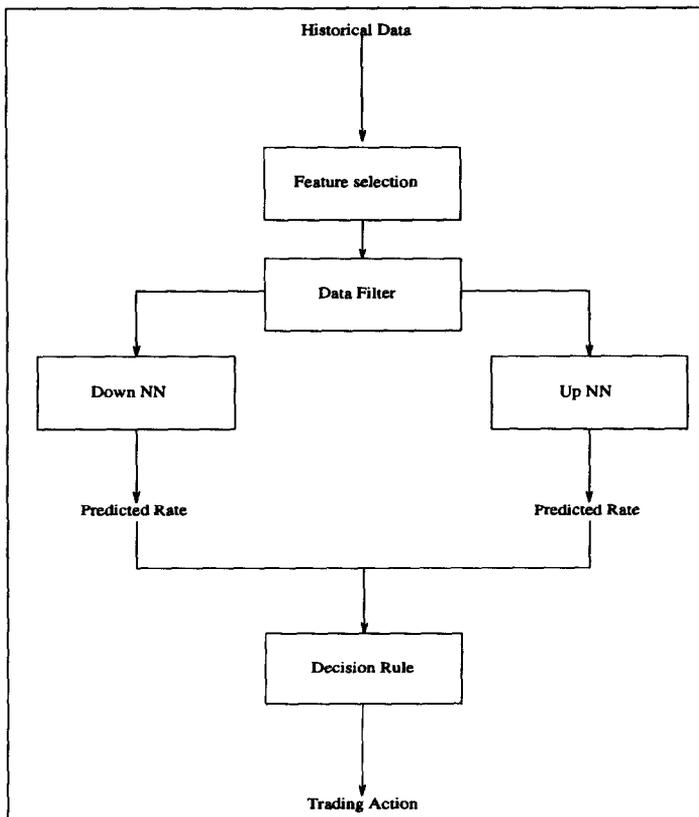

Fig. 1. System architecture.

that preserves as much information as possible. This issue is important because fewer features per pattern lead to faster computation and require less training patterns for successful generalization. Most feature selection processes rely on fundamentally sound statistically based techniques [11]. These techniques are practical, easy to understand, and easily implemented. However, they suffer from *instability* problems, meaning that small data perturbations lead to drastic changes in the final reduced feature set [3]. This problem is especially pronounced for stock market models because the data is non-stationary and very noisy. For these reasons, the feature selection procedure adopted in this paper is to use several selection techniques and criteria, then combine the partial results using a ranking process proposed in [7]. Techniques are discussed in Section 2.1.1, criteria in Section 2.1.2, and the ranking process in Section 2.1.3.

2.1.1. Selection techniques

Each feature selection technique is a search algorithm that attempts to determine a subset of the existing features which maximizes the differences between the classes based on some criteria. This section briefly describes the selection techniques used in the proposed feature selection process, all of which are described in more detail in [11].

The *Sequential Forward Search* selection technique is a greedy algorithm that begins with an empty feature set and adds features to it one at a time. The first feature added is the one deemed to be the best according to the selection criteria. The next feature added is the one which results in the largest improvement when considered in conjunction with the first feature. Similarly, the ith feature added is the one that results in the largest improvement when considered in conjunction with the previous $i - 1$ features. The *Sequential Backward Search* selection technique is similar to the sequential forward search, except that the initial set contains all the features, and features are removed from this set one at a time. The first feature removed is the one that results in the smallest degradation when the remaining features are considered together. This process repeats until the feature set reaches a predetermined size.

Although both the Sequential Forward Search and the Sequential Backward Search consider features in combination and as such are fine grained techniques, both require significant computing time. In addition, with Sequential Forward Search once a feature is added to the features set it cannot be removed. With Sequential Backward Search once a feature is removed it cannot be added later. As such, neither Sequential Forward Search nor Sequential Backward Search guarantee that an optimal set of features is achieved.

2.1.2. Selection criteria

Each feature selection criterion is based either on a measure of the distance between classes or an estimation of the classification error. Therefore the *selection criteria* objective is to either maximize the separation between classes using some distance measure of intra-class separation or to minimize the estimated classifica-

tion error. This Section describes the selection criteria used in the proposed feature selection process, all of which are described in more detail in [11].

The Euclidean, Patrick-Fisher, Mahalanobis, and Bhattacharyya distances used in this study are all means of measuring the multidimensional separation between classes. The *Euclidean* distance is measured as

$$\sqrt{(M_2 - M_1)^T (M_2 - M_1)}, \tag{1}$$

and the *Patrick-Fisher* distance as

$$\left(\frac{\Sigma_1 + \Sigma_2}{2}\right)^{-1} (M_2 - M_1), \tag{2}$$

where M_1 and M_2 are the mean vectors of class one and two respectively (e.g. M_1 is computed by averaging each feature's values for data in class one) and Σ_1 and Σ_2 are the corresponding covariance matrices. The *Mahalanobis* distance is measured as

$$(M_2 - M_1)^T \Sigma^{-1} (M_2 - M_1), \tag{3}$$

where the data from both classes is used to compute one covariance matrix Σ. And finally, the *Bhattacharyya* distance is measured as

$$\frac{1}{8} (M_2 - M_1)^T \left(\frac{\Sigma_1 + \Sigma_2}{2}\right)^{-1} (M_2 - M_1) + \frac{1}{2} \ln \frac{\left|\frac{\Sigma_1 + \Sigma_2}{2}\right|}{\sqrt{|\Sigma_1||\Sigma_2|}}. \tag{4}$$

Notice that the Bhattacharyya distance does not assume equal covariance matrices like the Mahalanobis distance. The first term in Eq. (4) measures the class separability due to the mean difference, while the second term measures the class separability due to the covariance-difference.

The final selection criterion used in this study is the *Estimated Minimal Error Probability* that estimates Bayes error for the data set by applying the K-nearest neighbor classifier [8] to the training set utilizing the leave one out approach. This criterion uses the Euclidean distance to determine the K training patterns closest to the selected pattern. The selected pattern is assigned to the class that the majority of the identified K training patterns belong to. This process is repeated for all patterns in the training set and the percentage of missclassifications computed, which becomes the estimate for Bayes error. This tends to overestimate the error and as such gives a very conservative error estimate. The selection criterion then becomes finding a set of features that minimizes the estimated Bayes error.

2.1.3. The ranking feature selection process

The approach proposed here for dealing with the instability of statistically based feature selection techniques for noisy and non-stationary data is to perform feature selection using several combinations of selection techniques and selection criteria

and to integrate the obtained partial results using a ranking process. The *ranking process* uses a specific technique and criterion combination to determine a rank ordering of the features from best to worst with scores assigned to the features based on this ordering (one is the best, p is the worst, where p is the total number of features). The process is repeated using combinations of other techniques and criteria explained earlier and the scores are summed. The reduced feature set is then comprised of the q features with the lowest scores.

For example, first the minimization of the Euclidean distance might be used as the selection criterion for the sequential forward search technique, which rank orders all the features from best to worst. A score is assigned to each feature based on this ranking (a score of 1 is the best), and the process repeated using the minimization of the Euclidean distance as the selection criterion for the sequential *backward* search technique, with the resulting score added to the previous score. The process continues using the forward and backward search techniques with other criteria (Bhattacharyya, Mahalanobis, Patrick-Fisher, and Estimated Minimal Error). The result of the process is a ranking of all features from best to worst according to the aggregated results of all eight combinations of selection techniques and criterion.

Had the initial data set been too large, this process would have been prohibitively expensive. For such cases, a two-stage ranking process is proposed in which the first phase reduces the initial feature set to an intermediate set of manageable size using coarser, computationally less demanding feature selection techniques. Then in the second phase the finer but more expensive Sequential Forward and Sequential Backward selection techniques are applied to the intermediate feature set to derive the final set [7].

2.2. Return rate prediction

The return rate prediction sub-system consists of a data filtering component and two NNs that are trained using the backpropagation algorithm and an on-line learning scheme (see Fig. 1). The filtering component's objective is to reduce the degree of noise in the data and then split the remaining patterns into two disjoint sets. The Up NN is trained on the set containing patterns with positive target returns and the Down NN on the set containing patterns with negative target returns. Once both NNs are trained, the test pattern is presented to each and the corresponding predictions are collected. A decision rule base is applied to these predictions and a buy/sell recommendation made as explained in Section 2.3.

The on-line learning scheme consists of a sequence of training/prediction sessions where the NNs are retrained after each session using more recent information. This is achieved by training the NNs using the backpropagation algorithm [15] and patterns from a fixed size window covering a continuous time segment of historic data. The target return for the time unit immediately following the window is predicted by both NNs and the predictions used by the rule base. Then the training window is shifted forward one time unit (i.e. one trading day or one month), the patterns from the new window used to retrain the NNs, and a

prediction made for the next time unit. This process is repeated until the data set is exhausted.

For each training session the target return corresponding to each pattern in the window is compared to a *filtering threshold* value h. If the return is greater than h the corresponding pattern is added to the Up NN training set, if the return is less than $-h$ the pattern is added to the Down NN training set. Any pattern with a target return between $-h$ and h is discarded.

For example, suppose that the training window size is m and that at time t the test pattern is d_t, which means that the training window contains patterns d_{t-m} through d_{t-1}. First, the patterns in the training window (d_{t-m} through d_{t-1}) are separated into Up NN and Down NN training sets using the threshold value h as described. Next, both NNs are trained using their respective training sets, and asked to predict the target return for the test pattern d_t. Once the predictions are collected and sent to the decision rule base, the training window is shifted forward one time unit so that the new test pattern is d_{t+1} and the new training window contains patterns d_{t-m+1} through d_t, and the process repeated. This continues until the end of the ordered data set is reached.

2.3. Decision rule base

The predicted returns from both NN components are used as input to the decision rule base component (see Fig. 1). This component analyzes the predicted returns and outputs a buy/sell recommendation that is used to establish either a long or short position in the market. A long position means purchasing an asset for later resale, while a short position means selling a borrowed asset now and purchasing it later.

This study examines three different decision rule bases. For each rule base the predicted market return from the Up NN (denoted by r_u) is compared to the predicted market return from the Down NN (denoted by r_d). Each rule base recommends a long position in the market if $r_u > 0$ and $r_d \geq 0$, and a short position if $r_u \leq 0$ and $r_d < 0$. Otherwise the rule base computes the normalized difference *diff*, which is a measure of the spread between r_u and r_d, normalized into a range from 0 to 1. It reflects the degree of confidence each NN has in its prediction with respect to the other. A large *diff* means that one of the two NNs is more confident in its prediction than the other, while a *diff* closer to zero means that each has about the same degree of confidence in its prediction. *Diff* is computed as

$$diff = \frac{\max\{|r_u|, |r_d|\} - \min\{|r_u|, |r_d|\}}{\max\{|r_u|, |r_d|\}}, \tag{5}$$

Each rule base compares *diff* to a predefined *decision threshold* value y, and determines a buy/sell recommendation as follows:
- **Rule base 1: Maintain current position until a clear buy/sell recommendation is received.**
  This rule base specifies that if the system is unsure as to what recommendation to make, the action is to do nothing and maintain the old position. Under

these rules, if $r_u \leq 0$ and $r_d \geq 0$ the system recommends maintaining the current position (i.e., do nothing). If $r_u > 0$, $r_d < 0$, and *diff* > *y* the rule base recommends a long position providing $r_u > |r_d|$, and a short position providing $r_u < |r_d|$. Otherwise *diff* ≤ *y* and the recommendation is to maintain the current market position.

- **Rule base 2: Hold a long position in the market unless a clear sell recommendation is received.**
  This rule base takes advantage of the common *a priori* knowledge that over the past 65 years the market has increased at an average annual rate greater than 10%. Stated another way, this means that given no other information the odds are that the market will increase. This is, in fact, the whole premise behind the buy and hold strategy. The difference between rule base two and rule base one is the actions taken under uncertainty. In this instance the action is to take a long position in the market. Under these rules, if $r_u > 0$, $r_d < 0$, *diff* > *y*, and $r_u < |r_d|$, the system recommends a short position. Otherwise the recommendation is to take a long position.

- **Rule base 3: Stay out of the market unless a clear buy / sell recommendation is received.**
  Again, the difference between rule base three and the previous rules is the action taken when the system is uncertain as to what recommendation to make. In case of uncertainty, the rule base three action is to exit the market. More precisely, if $r_u \leq 0$ and $r_d \geq 0$ the system recommends exiting the market (i.e. if the current position is long then sell, if it is short then buy). If $r_u > 0$, $r_d < 0$, and *diff* > *y* the system recommends a long position providing $r_u > |r_d|$, and a short position providing $r_u < |r_d|$. Otherwise *diff* ≤ *y* and the recommendation is to exit the market.

2.4. Performance measures

The most important criterion when measuring the performance of a stock market prediction model is whether it will make money and how much. Therefore the model's annual rate of return (*ARR*) is computed as

$$ARR = \frac{k}{n} \sum_{i=1}^{n} r_i, \tag{6}$$

where:
- k is the number of trading time units per year (i.e. 253 for daily trading, 12 for monthly trading);
- n is the total number of trading time units for the experiment (e.g. 2,530 for daily trading, 120 for monthly trading in an experiment lasting 10 years);
- r_i is the rate of return for time unit i.

The sum, $\sum_{i=1}^{n} r_i$, is computed by either adding, subtracting, or discarding the actual returns for the S & P 500 index. If the system recommends a long position, the actual return is added to the sum; if a short position is recommended, the

return is subtracted; or if the recommendation is to exit the market, the return is discarded.

It is also important to minimize transaction costs by controlling excessive trading (e.g. a 10% return with 50 trades is more profitable than a 10% return with 100 trades). Therefore the break even transaction cost (*BETC*), which may be viewed as the return per trade, is computed as

$$BETC = \frac{1}{s} \sum_{i=1}^{n} r_i,$$
(7)

where s is the total number of trading transactions, while r_i and n are defined as previously [12]. A *trade* is defined as any action that changes a market position. For example, exiting the market constitutes a single trade (i.e. a buy trade to cover a short position or a sell trade to cover a long position), while switching from a short position to a long position constitutes two trades (i.e. one buy trade to cover the short position and another buy to establish the long position).

3. Results and analysis

The system described in Section 2 is used for both daily and monthly S & P 500 stock market buy/sell recommendations. The daily historic data used in this experiment is 2,273 ordered financial time series patterns from the period January 1, 1985 to December 31, 1993. The first 1,000 patterns from January 1, 1985 to December 19, 1988 comprised the initial training window, whereas actual predictions were made for the 1,273 patterns from December 20, 1988 to December 31, 1993. Each pattern in the initial data set contained 24 monthly and 8 daily features as shown in the Appendix (Table 3). This was reduced to six features in the final data set, also shown in Table 3, by using the ranking feature selection process described in Section 2.

The monthly historic data consisted of an initial training window of 162 patterns formed using the data from January 1973 to February 1987 and actual predictions made for the 70 month period from March 1987 to December 1992. The initial monthly data set containing 29 monthly features was reduced to eight in the final data set by using the feature selection process described in Section 2. Both the initial and the reduced feature sets are shown in the Appendix (Table 4).

3.1. Daily results

The importance of the proposed feature selection and noise removal components were first tested using a *single NN system*. In this system the data was not split into up and down sets and a single NN replaced the dual NN component, making the decision rule base unnecessary. Results for a single NN system using the initial set of features, the reduced set, and the reduced set with a filtering threshold value of 0.5% are shown in Table 1 with the system parameters

Table 1
Results for the single NN system

	Daily			Monthly		
Preprocessing	ARR	BETC	Trades	ARR	BETC	Trades
Initial features	−2.16%	−0.00%	905	−1.67%	−0.00%	62
Reduced features	2.86%	0.00%	957	−3.33%	−0.01%	56
Reduced features and 0.5% noise removal	5.61%	0.01%	476	−2.97%	−0.01%	52

described in Table 2. All system parameter values were determined experimentally through the trial and error approach on a small subset of the training data (3 month period). It is important to note that the NN parameters where determined in the initial stage and fixed for the remainder of the experiment. To simplify the experimentation process, no effort was made to optimize these parameters dynamically as the training window shifted. In these experiments the *ARR* improves with feature reduction and noise removal. This provides evidence that in daily trading both feature selection and data filtering improve the NN predictive capability.

Several experiments with the dual NN system described in Section 2 were conducted using the reduced feature set from Table 3 and the system parameters from Table 2. Note that the training window size for the dual NN experiments is larger than for the single NN. This window size increase is needed since the training window for the dual NN system is split into 3 disjoint sets. The first set, consisting of all patterns with a target rate greater than a filtering threshold h, is used to train the Up NN. The second set, consisting of all patterns with a target rate less than $-h$, is used to train the 'Down NN'. Finally, the third set, consisting of all patterns with a target return between $-h$ and h, is discarded. Consequently, to ensure an adequately sized training set for both NN components in the dual NN

Table 2
System parameter values

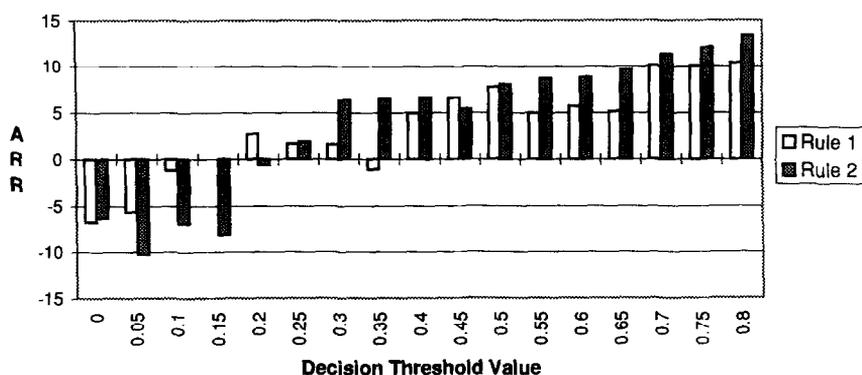| Parameter	Daily	Monthly
Topology (Single NN, initial features)	32-4-1	29-4-1
Topology (Single NN, reduced features)	6-4-1	8-3-1
Topology (dual NN)	6-4-1	8-3-1
Training window size (dual NN)	1000	162
Training window size (single NN)	250	162
Activation function	Tangent hyperbolic	Tangent hyperbolic
Learning rate	0.3	0.05
Tolerance	0.00001	0.00001
Number of iterations	5000	5000

Fig. 2. ARR comparison using daily data and filtering threshold $h = 0.5\%$.

system, it is necessary to have a larger window size. For the dual NN system, extensive experiments are conducted varying the filtering and decision thresholds h and y. The filtering threshold h is varied from 0.25% to 1.25% in increments of 0.25% with the best results obtained using a filtering threshold $h = 0.5\%$. The decision threshold y is varied from 0 to 0.80 in increments of 0.05. The upper boundary valued of 0.80 is used because larger values did not improve the results. A comparison between the *ARR* for rule one and rule two and the number of trades for rule one and rule two using a filtering threshold value $h = 0.5\%$ are shown in Figs. 2 and 3 respectively. The best annual rate of return was 13.35% and was obtained using rule base two with thresholds $h = 0.5\%$ and $y = 0.80$. In comparison, the annual rate of return for the same period using the buy and hold strategy is 11.23% and the best return for the single NN is only 5.61%. Results for decision rule three are significantly less than the results using decision rules one and two (best *ARR* and *BETC* are 6.50% and 0.08% respectively) and as such are not presented.
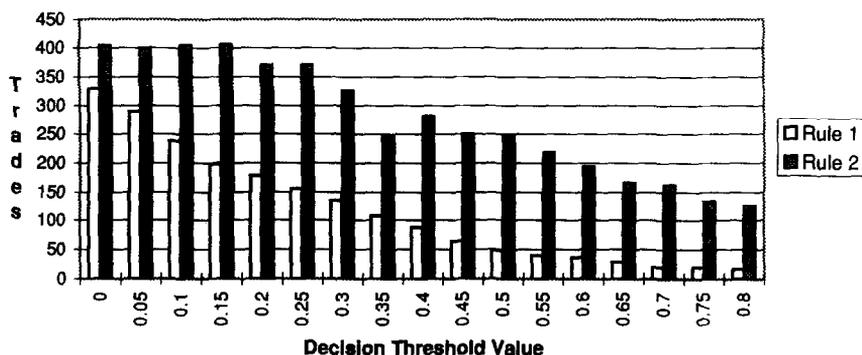

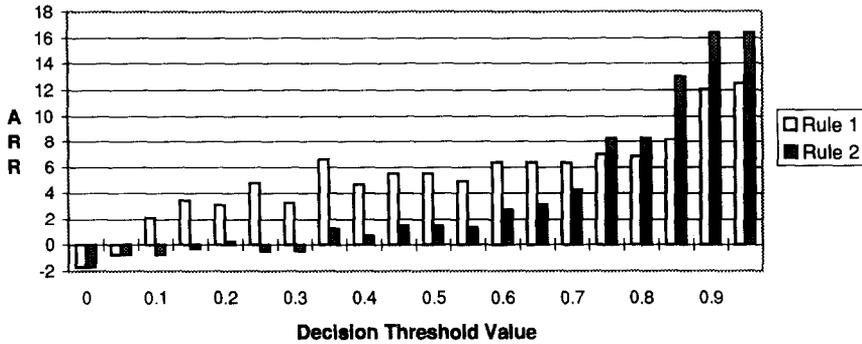Fig. 3. Number of trades comparison using daily data and filtering threshold $h = 0.5\%$.

Fig. 4. ARR comparison using monthly data and filtering threshold $h = 0.0\%$.

3.2. Monthly results

Results for a single NN system using the initial set of features, the reduced set, and the reduced set with a filtering threshold value of 0.5% are also shown in Table 1 while the system parameters for both single NNs are shown in Table 2. Similar to the daily experiments, system parameter values were determined experimentally on a small subset of the training data with no effort made to optimize these values for each training window. The reduced feature set gave results that were worse than those achieved using the complete feature set, but without drastic information loss. In addition, noise removal did not significantly improve the results, indicating that noise removal may not be necessary for the smoother monthly data.

Several experiments with the dual NN system described in Section 2 were conducted using the reduced feature set from the Appendix (Table 4) and the system parameters from Table 2. For the dual NN system, extensive experiments are conducted varying the thresholds h and y. The filtering threshold h is varied from 0.0% to 3.00%. Observe that for daily experiments a corresponding upper
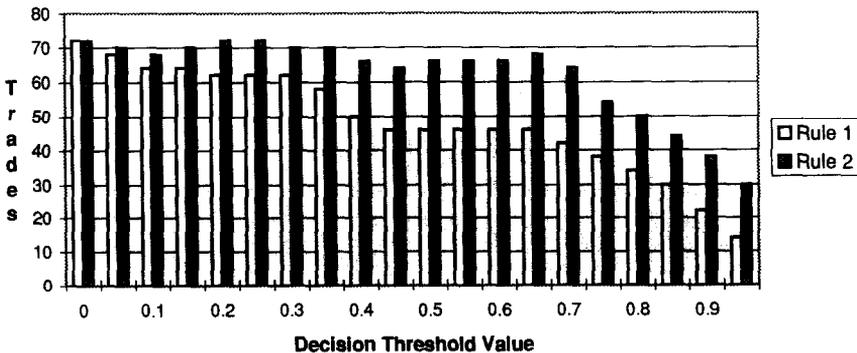

Fig. 5. Number of trades comparison using monthly data and filtering threshold $h = 0.0\%$.

bound is 1.25% versus the 3.00% used here. This 3.00% study interval is not feasible for daily data since it would remove too much data. For all rules the best results were achieved using a filtering value $h = 0\%$, meaning that there was no need for noise reduction as earlier indicated by the single NN results. The decision threshold y is varied from 0 to 0.95 in increments of 0.05. The upper boundary value of 0.95 is used because larger values did not improve the results. A comparison between the ARR for rule one and rule two and the number of trades for rule one and rule two are shown in Figs. 4 and 5 respectively. The best annual rate of return was 16.39% and was obtained using rule base two with thresholds $h = 0\%$ and $y = 0.95$. In comparison, the annual rate of return for the buy and hold strategy was 8.76% and the best return for the single NN was -1.67%. Results for decision rule three are significantly less than the results using decision rules one and two (best ARR and BETC are 8.76% and 1.51% respectively) and as such are not presented.

4. Conclusions and future research

The system proposed in this paper is comprised of a preprocessing component for feature selection, a filtering component for noise removal and pattern separation, two specialized NN components for return predictions, and a decision rule component for buy/sell recommendation. Various experiments using this system to predict S & P 500 index movements were conducted and associated annual rates of return and returns per transaction computed. Although the obtained results are promising, it is important to note that no claims are made that the proposed system is better than the buy and hold strategy when considering transaction costs.

Comparison of the results achieved by the dual-NN system to that of the single NN shows that the dual NN system gives a larger return with fewer trades. In addition, dual NN experiments with the appropriately selected filtering and decision thresholds managed to achieve an annual rate of return greater than that of the buy and hold strategy. Comparing the daily to the monthly experimental results, it seems evident that the proposed system yields better returns when using slower data sampling (e.g. monthly rather than daily data).

The proposed system is still in development and research in progress might lead to further improvements. For instance, sampling rates other than those used in this study are also possible (e.g. inter-day trading, quarterly trading). Further research is needed to determine the optimal sampling rate for the proposed trading system. In addition, no attempt was made to optimize the NN parameters dynamically as the training window shifted. It is possible that optimized learning parameters such as the learning rate may lead to better results. It may also be possible to improve system performance through optimization of the system specific filtering threshold parameter. In this study, several experiments utilizing different values for the filtering threshold parameter were performed, but no effort was made to determine the optimal value for this parameter as this would be computationally prohibitively expensive. Each new filtering threshold value determines a different partitioning of the data set requiring extensive retraining of both NNs.

This study did attempt to optimize the decision threshold parameter as new values of the parameter did not require NN retraining and as such this optimization was computationally less expensive. However, the current decision rule bases are fairly simplistic. Possible improvements might be obtained by incorporating technical information like moving averages and exponential averages into the system.

It also might be possible to improve results by further restricting the scope of each NN by incorporating additional NNs. For example, in a four network system the first NN can be trained on large down movements, the second NN on small down movements, the third on small up movements, and the fourth on large up movements. Further improvements might also be obtained by incorporating prior knowledge and constructive NN learning [10], or a recurrent network topology [4]. Experiments using additional NNs, an expert system, and combinations of the two to analyze the existing system information and determine a market direction are in progress.

Acknowledgments

Appendix

Table 3
Features for daily experiments

Initial Features	Reduced
Return on 30 year Government Bonds	X
30 year Government Bond index	
Rate of change in the return on Government Bonds	
Rate of change in the return on Government Bonds lagged 1 month	
Rate of change in the return on Government Bonds lagged 2 months	
Rate of change in the return on Government Bonds lagged 3 months	
Rate of change in the return on Government Bonds lagged 4 months	
Rate of change in the return on Government Bonds lagged 5 months	
Return on U.S. Treasury Bills	
U.S Treasury Bill Index	
Rate of change in the return on T-bills	
Rate of change in the return on U.S Treasury bills lagged 1 month	
Rate of change in the return on U.S Treasury bills lagged 2 months	X
Rate of change in the return on U.S Treasury bills lagged 3 months	X
Rate of change in the return on U.S Treasury bills lagged 4 months	
Rate of change in the return on U.S Treasury bills lagged 5 months	
The CPI	

The percentage increase in the CPI
The rate of change in the percentage increase in the CPI
The rate of change in the percentage increase in the CPI lagged 1 month
The rate of change in the percentage increase in the CPI lagged 2 months
The rate of change in the percentage increase in the CPI lagged 3 months
The rate of change in the percentage increase in the CPI lagged 4 months
The rate of change in the percentage increase in the CPI lagged 5 months
The S&P Composite Index
The S&P Composite Index lagged 1 day
The S&P Composite Index lagged 2 days
The return on the S&P Composite Index   X
The return on the S&P Composite Index lagged 1 day   X
The return on the S&P Composite Index lagged 2 days   X
The return on the S&P Composite Index lagged 3 days
The return on the S&P Composite Index lagged 4 days

Table 4
Features for monthly experiments

Initial Features	Reduced
Return on 30 year Government Bonds	
30 year Government Bond index	X
Rate of change in the return on Government Bonds	
Rate of change in the return on Government Bonds lagged 1 month	
Rate of change in the return on Government Bonds lagged 2 months	
Rate of change in the return on Government Bonds lagged 3 months	X
Rate of change in the return on Government Bonds lagged 4 months	X
Rate of change in the return on Government Bonds lagged 5 months	X
Rate of change in the return on Government Bonds lagged 6 months	
Return on U.S. Treasury Bills	
U.S Treasury Bill Index	X
Rate of change in the return on T-bills	
Rate of change in the return on U.S Treasury bills lagged 1 month	X
Rate of change in the return on U.S Treasury bills lagged 2 months	
Rate of change in the return on U.S Treasury bills lagged 3 months	
Rate of change in the return on U.S Treasury bills lagged 4 months	
Rate of change in the return on U.S Treasury bills lagged 5 months	
Rate of change in the return on U.S Treasury bills lagged 6 months	
The CPI	X
The percentage increase in the CPI	
The rate of change in the percentage increase in the CPI	
The rate of change in the percentage increase in the CPI lagged 1 month	
The rate of change in the percentage increase in the CPI lagged 2 months	
The rate of change in the percentage increase in the CPI lagged 3 months	
The rate of change in the percentage increase in the CPI lagged 4 months	
The rate of change in the percentage increase in the CPI lagged 5 months	
The rate of change in the percentage increase in the CPI lagged 6 months	
The S&P Composite Index	X
The return on the S&P Composite Index	

References

[1] Y.S. Abu-Mostafa, Financial market applications of learning from hints, in A.N. Refenes, ed., *Neural Networks in the Capital Markets* (Willey, England, 1994).

[2] F. Black and M. Scholes, The pricing of options and corporate liabilities, *J. Political Economy* 81 (May-June 1973).

[3] L. Breiman, The heuristics of instability in model selection, Technical Report No. 416, Statistics Department, University of California, Berkeley, CA, 1994.

[4] A. Burgess and D. Bunn, The use of error feedback terms in neural network modeling of financial time series, *Proc. 1994 Neural Networks in the Capital Markets Conf.* Pasadena, CA (1994).

[5] T. Chenoweth and Z. Obradovic, Feature selection for predictive models of the stock market, *Proc. 1994 Neural Networks in the Capital Markets Conf.*, Pasadena, CA (1994).

[6] T. Chenoweth and Z. Obradovic, A multi-component approach to stock market predictions, *Proc. Third Int. Conf. on Artificial Intelligence Applications on Wall Street*, New York, NY (1995) 74–79.

[7] T. Chenoweth and Z. Obradovic, An explicit feature selection strategy for predictive models of the S & P 500 Index, *NeuroVest^{TM} J.* 3 (6) (1995) 14–21.

[8] T.M. Cover and P.E. Hart, Nearest neighbor pattern classification, *IEEE Trans. Inform. Theory* IT-13 (1967) 21–27.

[9] G. Cybenko, Approximation by superpositions of a sigmoidal function, *Math. Control, Signals, and Systems* 2 (1989) 303–314.

[10] J. Fletcher and Z. Obradovic, Combining prior symbolic knowledge and constructive neural networks, *Connection Science: J. Neural Computing, Artificial Intelligence and Cognitive Research* 5 (3–4) (1993) 365–375.

[11] K. Fukunaga, *Introduction to Statistical Pattern Recognition* (Academic Press, San Diego, CA, 1990).

[12] J. Hutchinson, A radial basis function approach to financial time series analysis, PhD Thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, 1993.

[13] M. Jensen, Some anomalous evidence regarding market efficiency, *J. Financial Economics* 6 (1978) 95–101.

[14] J. Markowitz, *Portfolio Selection: Efficient Diversification of Investments* (John Wiley & Sons, New York, NY, 1959).

[15] Rumelhart et al. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, 1 and 2 (MIT Press, Cambridge, MA, 1986).

[16] W. Sharpe, Capital asset prices: A theory of market equilibrium, *J. Finance* (Sep. 1964).

Tim Chenoweth (tchenowe@eecs.wsu.edu) received a B.S. degree in Mathematics in 1981 from the Coast Guard Academy, a M.B.A. in Finance from Washington State University in 1991, and is currently completing a M.S. in Computer Science and an Individual Interdisciplinary Ph.D. combining Business and Computer Science, both from Washington State University. He was an active duty officer in the Coast Guard from 1981 to 1989. The objective of his current research is to use advanced technologies to model financial markets.

Zoran Obradovic (zoran@eecs.wsu.edu) received the B.S. degree in Applied Mathematics, Information and Computer Sciences in 1985; the M.S. degree in Mathematics and Computer Science in 1987, both from the University of Belgrade; and the Ph.D. degree in Computer Science from the Pennsylvania State University in 1991. He was a systems programmer at the Department for Computer Design at the Vinca Institute, Belgrade, from 1984 to 1986, and has been a research scientist at the Mathematical Institute of the Serbian Academy of Sciences and Arts, Belgrade, since then. At present, he is an Assistant Professor in the School of Electrical Engineering and Computer Science, Washington State University, Pullman, WA 99164-2752, USA. The objective of his current research is to explore applicability of neural networks technology to large scale classification and time series prediction problems in very noisy domains.