A NEURAL NETWORK APPROACH FOR THE DIRECTIONAL PREDICTION OF A STOCK MARKET: AN APPLICATION TO THE AUSTRALIAN ALL ORDINARY INDEX

C. D. Tilakaratne
Department of Statistics, University of Colombo, Sri Lanka
cdt@stat.cmb.ac.lk

ABSTRACT

This study investigates the capability of neural network-based approaches for predicting the direction (up or down) of a stock market. The Australian Stock market index (AORD) was selected as the target market. It includes several aspects: univariate feature selection from the historical time series of the target market, intermarket analysis for finding the most relevant influential markets, investigations of the effect of time cycles on the target market and the discovery of the optimal neural network architectures. It was found that the relative return series of the Open, High, Low and Close prices of the target market, show 6-day cycles during the study period of about fourteen years. Multilayer feedforward neural networks trained with a backpropagation algorithm were used for the experiments. The best neural network developed in this study has achieved 87%, 81% 83% and 81% accuracy respectively in predicting the next-day direction of the relative return of the Open, High, Low and Close prices of the AORD. The architecture of this network consists of 33 input features, one hidden layer with three neurons and four output neurons. The best input features set includes the relative returns from one to six days in the past of the Open, High, Low and Close prices of the target market, the day of the week, and the previous day's relative return of the Close prices of the US S&P 500 Index, US Dow Jones Industrial Average Index, US Gold/Silver Index, and the US Oil Index.

Keywords: Neural network, Directional prediction, Australian stock market, Feature selection

1. INTRODUCTION

Before the 1980s, attempts to model financial market data in order to predict future market directions were unsuccessful due to the inherent complexity of the data. The efficient market hypothesis claims that financial markets are a random time series and, therefore unpredictable on the basis of any amount of publicly available knowledge (Malkiel, 1985).

Until the late 1980s, most quantitative approaches to testing this hypothesis were based on linear time series modelling (White, 1988). Chenoweth *et al*., (1996) stated that it is very hard to find statistically significant market inefficiencies using standard linear time series modelling, since such linear approaches are not capable of identifying dynamic or non-linear relationships in financial data. Furthermore, it is widely accepted classical statistical methods, such as multiple linear regression and ARIMA models are limited to applications with non-stationarity and non-linearities in the stock market data (Refenes *et al*., 1994; Kolarik and Rudofer, 1994).

Weiss and Kulikowski (1991) suggested that an appropriate nonparametric machine-learning technique might be able to discover more complex non-linear relationships through supervised learning from examples. Such new approaches to financial modelling have been developed during the last two decades.

Cybenko (1989) found that Neural Networks (NNs) are powerful computational systems that can approximate almost any non-linear continuous function on a compact domain to any desired degree of accuracy. NNs have the ability to learn non-linear chaotic systems, without explicit assumptions on prior distributions. Therefore, in general, they are preferred to other statistical techniques of forecasting stock markets.

White (1988) used NNs to predict IBM daily stock returns. Since then, a great deal of research has been carried out in the use of NNs for predicting stock market behaviour. However, most of this research has been targeted at predicting the US and other international stock markets (Yao and Poh, 1995; Chenoweth and Obradovic, 1995; Gately, 1996; Mizuno, 1998; Qi, 1999; Olson and Mossman, 2002; Egeli, *et al.* , 2003) and no research has been published on predicting the direction of the Australian stock market, before our work (Pan *et al*., 2003; Tilakaratne *et al*., 2006). Every stock market is different, and has its unique 'personality' and position in the international economic systems (Pan *et al*., 2003). Therefore, prediction related to the Australian stock market, is an important and unique issue for the Australian finance academy and industry and also poses a great challenge to artificial intelligence and statistics.

Application of NN techniques to predict the financial markets is one of the biggest challenges in artificial intelligence, since late 80's. Human intelligent for predicting financial markets may well be inappropriate. Therefore, developing an artificial intelligence system for predicting financial markets is not simply a matter of re-engineering human expert knowledge, but rather iterative process of knowledge discovery and system improvement through data mining, knowledge engineering, theoretical and data-driven modelling, as well as trial and error experimentation (Pan et al., 2003). Therefore, the application of neural network techniques to predict any financial market is a challengeable task.

A number of previous studies have attempted to predict the price levels of stock market indices (Egeli *et al*., 2003; Gencay and Stengos, 1998; Qi, 1999; Safer, 2003). However, in the last few decades, there have been a growing number of studies attempting to predict the direction or the trend movements of financial market indices (Cao and Tay, 2001; Huang *et al*., 2005; Kim and Chun, 1998, Qi and Maddala, 1999; Wu and Zhang, 1997). Some studies have suggested that trading strategies guided by forecasts on the direction of price change may be more effective and may lead to higher profits (Wu and Zhang,, 1997). Leung *et al*. (2000) also found that the classification models based on the direction of stock return outperform those based on the level of stock return in terms of both predictability and profitability

The main objective of this study was to predict the direction (up or down) of the Open, High, Low and Close prices of the Australian All Ordinaries Index (AORD), one day into the future, with high reliability and accuracy, using artificial neural networks. The Australian stock market is one of the leading stock markets in the Asia Pacific region as well as being a major international market. The Australian economy has been one of the healthiest and most reliable in the world. Therefore, the Australian Stock Market has attracted a large number of international investors and traders. A method that could predict movements of the target market with high reliability and accuracy would be significantly useful for traders as well as policy makers.

The organisation of the paper is as follows: the next section describes the architecture of NNs, followed by the third section which explains how to train NNs and the evaluation of their performance. Section 4 focuses the feature selection. Section 5 describes the NN experiments and also presents the results. The final section gives the conclusions.

2. NN ARCHITECTURE

An artificial neural network is a computer program or hardwired machine that is designed to learn in a manner similar to human brain. Several different artificial neural network architectures have been developed over the past two decades. Common to all these architectures is the linking together of many neurons, or nodes, into numerous interconnected pathways of inputs and outputs. This research was limited to multiple feedforward neural networks.

2.1. FEEDFORWARD NN

Figure 1 depicts an example of a multi-layer feedforward neural network. A multi-layer feedforward neural network can have any number of layers and any number of units (neurons) per layer. The first layer is called the Input Layer and the last layer is called the Output Layer. The middle layers are called Hidden Layers. The network shown below has 4 neurons (or units) in the input layer, three neurons in the hidden layer, and one neuron in the output layer.

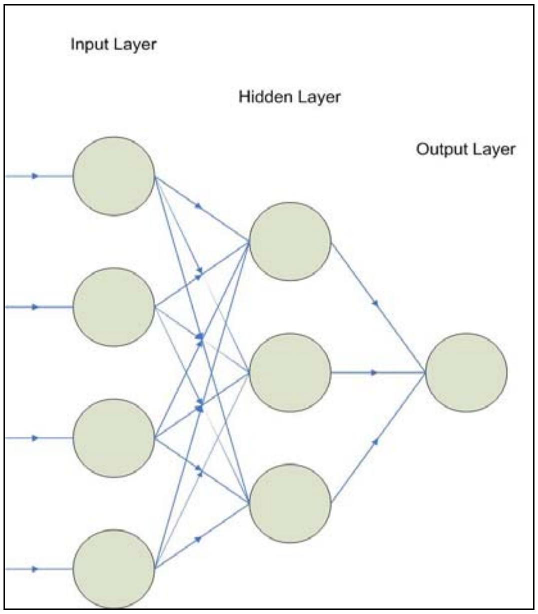

Figure 1: A three-layer feedforward NN

Each neuron-to-neuron connection is modified by a weight (or connection strength). In addition, each neuron has an extra input that is assumed to have a constant value of one, and the weight that modifies this extra input is called the bias. All the information propagates along the connections in the direction of network inputs to network outputs, hence the term feedforward.

The selection of the number of input neurons is usually done by feature selection. Once the input variables are selected the number of input neurons is simply the number of input features selected. The number of output neurons equals the number of predictions that the network is designed for.

Determination of the number of neurons per hidden layer is the most difficult and important task in the design of a multi-layer feedforward neural network. There is no specific formula or procedure to decide the number of hidden neurons. Selecting the optimal number of hidden neurons for a given input-output set-up, involves experimentation using real data.

In practice, neural networks with one and occasionally two hidden layers are widely used and have performed well enough. Increasing the number of hidden layers leads to increases in computer time, as well as the risk of instability and over-fitting (Kaastra and Boyd, 1996).

The input neurons simply pass on the input vector x (input variables). The neurons in the hidden layer(s) and the output layer are processing units. Each processing neuron is associated with an activation function.

The following equation gives the net input to the j^{th} neuron of the current hidden layer:

$$net_j = \sum_i w_{ij} l_i + b_j \tag{1}$$

where l_i is the output of the i^{th} neuron of the input layer or the previous hidden layer, w_{ij} is the weight of the link connecting i^{th} neuron of the previous layer to j^{th} neuron of the current hidden layer, and b_j is the bias associated with the j^{th} neuron of the current hidden layer.

The output of the j^{th} neuron of the current hidden layer is given by:

$$out_j = f(net_j) \tag{2}$$

where f is the activation function associated with the hidden neuron. For this study, the sigmoid function was used. This is the activation function most commonly used to train feedforward neural networks. Furthermore, if a network is to learn average behaviour, a sigmoid transfer function is suitable (Kaastra and Boyd, 1995). According to the objectives of this research, it is expected that the networks will learn average behaviour. This is another reason to use this activation function.

The input to and the output of the j^{th} neuron in the output layer can be defined similarly. Equations (3) and (4) give the input to and the output of the j^{th} neuron of the output layer respectively:

$$net'_j = \sum_i w'_{ij} l'_i + b'_j \tag{3}$$

where l'_i is the output of the i^{th} neuron of the previous hidden layer, w'_{ij} is the weight of the link connecting i^{th} neuron of the previous layer to j^{th} neuron of the output layer, and b'_j is the bias associated with the j^{th} neuron of the output layer.

$$out'_j = f'(net'_j) \tag{4}$$

where f' is the activation function associated with the output neuron. Usually, $f'(x) = x$ is sufficient.

3. NN TRAINING AND PERFORMANCE EVALUATION

In the problem of stock market prediction, the network is fed inputs as well as outputs (training data or sample data). Then the network learns the mapping from inputs to corresponding outputs. This is called Supervised Learning. In order for the network to learn the patterns of the data (that is to estimate the weights associated with connections of the network), a learning algorithm is needed. Backpropagation is the learning algorithm most commonly used for the feedforward neural networks. Associated with the training algorithm, there are two parameters: learning rate and momentum.

The learning rate controls how quickly and how finely a network converges to a particular solution. A large value for learning rate will lead to rapid learning but the weight (associated with connections of NN) may then oscillate, while low values imply slow learning. The proper value for learning rate depends on the application. Usually this value changes from 0 to 1 (Mehrotra, *et al*., 1997).

Momentum is the weight change required at a time *t*. A value close to 0 indicates that the past history does not have much effect on the weight change. However, a value close to 1 suggests that current error has little effect on the weight change (Mehrotra, *et al*., 1997).

A NN is considered to be efficient in prediction, only if it can predict the out of sample (test) data[1] well. A network with a large number of neurons may predict the sample data well, by over-fitting the input patterns, but may poorly predict the out of sample data.

A network is said to be well generalised when the input-output mapping computed by the network is approximately correct for the test data. In backpropagation learning, it is expected that a network become well trained. This means it has learned enough about the past to generalise to the future. A network that produces a high level of forecasting errors on the test set, but a low level of errors on training set, is said to overfit to the training data. This happens when testing of network is done after completing the training.

The size of the training data set also affects the prediction capability of a network. If the size of the training data set is too small, the network tends to overfit the data. The standard is to allocate the majority of data (for example 90%) into a training set and the balance into a test set. Sometimes the training set is further divided into a validation set and a training set. For instance, 70% of data are used for training, 20% of data are used for validation and the balance is used for testing. The training set is used to build the model while the validation set is used to test the training as it proceeds to determine the best point for early stopping or the best network architecture (number of neurons in the hidden layer). The test set is only used for the final estimation of the generalisation performance.

The performance of the neural networks can be evaluated by using several statistics. Since the interest is to predict the direction of the time series, the performance of the test data set can be evaluated by the Sign Correctness Percentage (SCP), which is given by the Equation (5) (Pan *et al*., 2003). The SCP indicates the percentage of predicted outputs with correct direction. In predicting a stock market, the most important issue is to predict the direction of the market rather than the value. The higher the value of the SCP, the better the performance of the neural network.

$$SCP = \frac{\left|\{sign(z_r) = sign(o_r) \mid r = 1, 2, ..., N_2\}\right|}{N_2} \quad (5)$$

where N_2 is the number of samples in the test data set.

4. FEATURE SELECTION ADD DATA PREPROCESSING

The features extracted from the target market (the AORD) and the intermarkets with some other features such as day of the week can be included in the input vector. According to the author's empirical knowledge the following markets were considered:

- US S&P 500 Index (GSPC)
- US NAS/NMS Composite Index (IXIC)

[1] Unseen data.

- US Dow Jones Industrial Average (DJI)
- US Gold/Silver Index (XAU)
- US Oil Index (XOI)

Daily relative returns[2] ($r(t)$) of the prices of the above markets and the AORD were considered. The study period consists of fourteen years of data (from 2nd January, 1990 to 20th February, 2004) from each market. This period includes both bullish and bearish market periods as well as market crashes (for instance crash due to September 11th attack). If no trading took place on a particular day, the rate of change of price should be zero. Therefore, before calculating the relative returns, the missing values of a price were replaced by the respective price of the last trading day.

Figures 2 to 5 depict the collelograms (plots of autocorrelation coefficients versus lag) of the relative returns of the Open, High, Low, and the Close prices of the AORD during the study period.

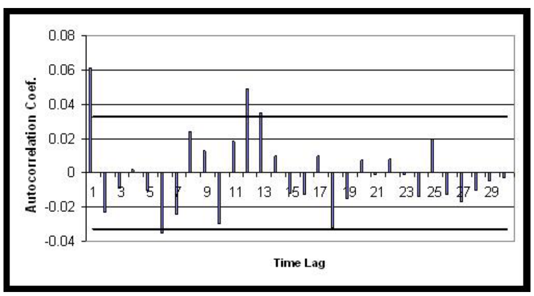

Figure 2: Correlogram of relative Returns of the Open price of the AORD

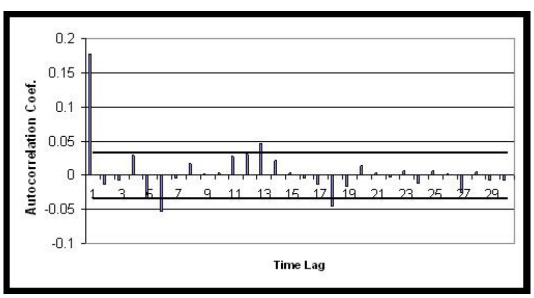

Figure 3: Correlogram of relative returns of the High price of the AORD

[2] $r(t) = \dfrac{P(t) - P(t-1)}{P(t-1)}$ where $P(t)$ is the price on day t. Relative returns are preferred to original prices since they facilitate the comparison of different stocks markets on an equal basis

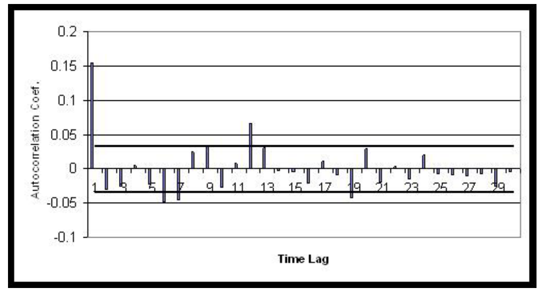

Figure 4: Correlogram of relative returns of the Low price of the AORD

Figure 2 shows the behaviour of relative returns of the Open price of the target market. It has spikes at lags 1, 6, 12, 13, and 18. The autocorrelation coefficients of larger lags are not significant. Therefore it can be concluded that this series is stationary. Also it shows an approximate 6-day cycle.

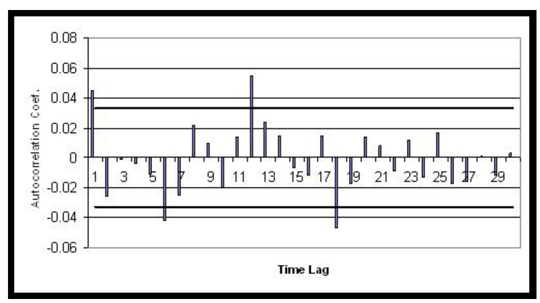

Figure 5: Correlogram of relative returns of the Close price of the AORD

Figure 3 shows spikes at lags 1, 6, 13, and 18. For higher lags the distribution of the autocorrelation coefficient becomes flat. This series also shows an approximate 6-day cycle.

The next figure (Figure 4) also shows the spikes at lags 1, 6, 7, 12, and 18. This behaviour indicates an approximate 6-day cycle in the considered series.

Figure 5 shows only a few spikes at lower lags (1, 6, 12, and 18), while the distribution of the autocorrelation coefficients becomes flat at higher lags. The important characteristic of the behaviour of these autocorrelation coefficients is that they show spikes at every 6th day, until 18th day. Therefore, it can be concluded that relative returns of the Close price of the AORD also show 6-day cycles.

The correlograms (plots of autocorrelation functions) of the relative return series of all prices (Open, High, Low and Close) of the AORD show an approximate 6-day cycle during the period considered. Hence it was decided to take relative returns of one to six days in the past into account when predicting the direction of movement of the prices of the AORD.

In order to find the influences of the considered inter-markets, the Pearson correlation coefficient and the partial correlation coefficient were used (see Table 01). It was found that the Close prices of the US S&P 500, the US Dow Jones Industrial Average, the US Gold/Silver Index and the US Oil Index affect the prices of the AORD.

Table 01: Pearson correlation coefficient between relative returns (at time t+1) of the Open price of the AORD and relative returns (at time t) of the Close prices of the Intermarkets

Intermarket	Correlation Coefficient
US S&P 500 Index	0.422**
US NAS/NMS Composite Index	-0.001NS
US Dow Jones Industrial Average	0.183**
US Gold/Silver Index	0.086**
US Oil Index	0.074**

**-Correlation is significant at the 0.01 level.

NS-Not Significant

Partial correlation coefficients confirmed that the effect of each market is significant even at the presence of other market.

The influence of the day of the week on the different prices of the target market was tested using the chi-square tests. The results revealed that the day of the week only affects the behaviour of the relative returns of the High price (since this is the only case where the test statistic is significant). Therefore, it is useful to include the day of the week as an input feature when predicting the direction of relative returns (at time t) of the High price for the AORD.

Finally, 33 variables were selected as input features to predict the direction of relative returns of the current day's (day t) of the four prices of the AORD simultaneously. They are:

- relative returns from 1 to 6 days in the past of the Open price of the AORD;
- relative returns from 1 to 6 days in the past of the High price the AORD;
- relative returns from 1 to 6 days in the past of the Low price the AORD;
- relative returns from 1 to 6 days in the past of the Close price of the AORD;
- previous day's relative returns of the Close prices of the US S&P 500 Index;
- previous day's relative returns of the Close prices of the DJI;
- previous day's relative returns of the Close prices of the XAU;
- previous day's relative returns of the Close prices of the XOI;
- 5 dummy variables (for instance Monday is presented as 1, 0, 0, 0, 0) to represent the day of the week.

5. NN EXPERIMRNTS

An attempt was made to find a suitable neural network model to predict the direction of the current days (day t) relative returns (that is positive or negative) of the Open, High, Low and

Table 2: Input and output variables

Input variables	$O(t\text{-}1)$ $O(t\text{-}2)$ $O(t\text{-}3)$ $O(t\text{-}4)$ $O(t\text{-}5)$ $O(t\text{-}6)$ $H(t\text{-}1)$ $H(t\text{-}2)$ $H(t\text{-}3)$ $H(t\text{-}4)$ $H(t\text{-}5)$ $H(t\text{-}6)$ $L(t\text{-}1)$ $L(t\text{-}2)$ $L(t\text{-}3)$ $L(t\text{-}4)$ $L(t\text{-}5)$ $L(t\text{-}6)$ $C(t\text{-}1)$ $C(t\text{-}2)$ $C(t\text{-}3)$ $C(t\text{-}4)$ $C(t\text{-}5)$ $C(t\text{-}6)$ $Y_1(t\text{-}1)$ $Y_2(t\text{-}1)$ $Y_3(t\text{-}1)$ $Y_4(t\text{-}1)$ $V_1(t)$ $V_1(2)$ $V_1(3)$ $V_1(4)$ $V_5(t)$
Output variables	$O(t)$ $H(t)$ $L(t)$ $C(t)$

the Close prices of the target market simultaneously using 33 input features that were identified (see Section 4). These networks consisted of 33 input neurons and 4 output neurons. Table 2 presents the variables correspond to the input and output layers.

- *O, H*, *L*, *C* - Relative returns of the Open, High, Low and Close prices of the AORD market respectively
- Y_1 - Relative returns of the Close price of the US S&P 500 Index
- Y_2 - Relative returns of the Close price of the US Dow Jones Industrial Average Index
- Y_3 - Relative returns of the Close price of the US Gold/Silver Index
- Y_4 - Relative returns of the Close price of the US Oil Index
- $V_1 \ldots V_5$ – Five dummy variables represent the 5 days of the week
- *ZO*, *ZH*, *ZL*, *ZC* – Predicted relative returns of the Open, High, Low and Close prices of the target market
- *t* – Current time

The most recent 10% of data was allocated for testing while the next most recent 20% of data (728 days) was set apart for validation; the remaining data (2549 days) was used for training. Three-layer feedforward neural networks with one hidden layer were trained with the 2549 records (data from 2nd January 1990 to 24th November 1999) allocated for training. These networks consisted of 33 input neurons and four output neurons. The number of neurons in the hidden layer was varied, as were the learning rate and the momentum coefficient.

All the trained neural networks were used to predict the direction of relative returns of all four prices of the AORD, from 24th September 2002 to 20th February 2004 (the test data set). The values of the SCP of the validation and test data were plotted against the number of neurons in the hidden layer (see Figure 6).

Although, there is a significant difference between the two values of the SCP when the number of hidden neurons equals 1, in all the other cases there is no significant difference between the respective values of the SCP. The highest value (83.05%) for the SCP for the test data set occurs when the hidden layer consists of three neurons. For the test data, this NN gave 81.36% accuracy in predicting the direction of the relative returns of the Close price. The accuracy values for the Open, High, and Low prices are 87.29%, 80.51%, and 83.06% respectively. Therefore, it is reasonable to state that a three-layer feedforward neural network with three neurons in the hidden layer is the best architecture for predicting the direction of relative returns of all four prices simultaneously.

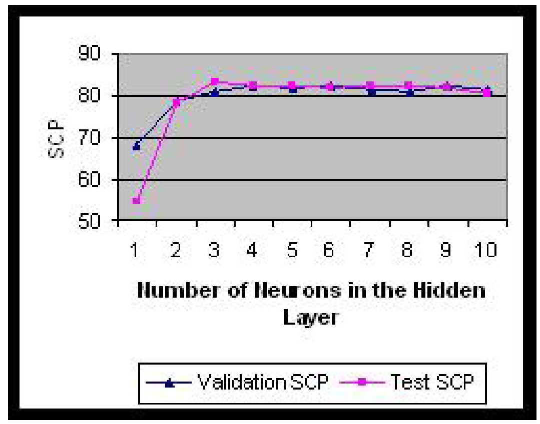

Figure 6: SCP of validation and test sets for different number of hidden neurons

5.1. COMPARISON OF PREDICTIONS WITH ACTUACL VALUES

Figures 7 to 10 compare the predicted values of the Open, High, and Low prices of the target market with their corresponding actual values.

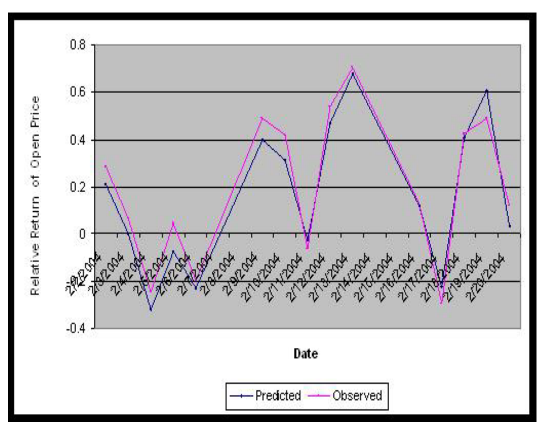

Figure 7: Predicted and actual values for the relative returns of the Open price versus date

According to the Figure 07, the network predicts the direction of the relative returns of the Open price of the target market for all days of the considered period correctly, except for 13th and 19th of February. Therefore, it can be stated that this neural network has the ability to predict the direction of the Open price of the target market accurately. The difference between the magnitude of the predicted and the actual values is also not significant in many cases.

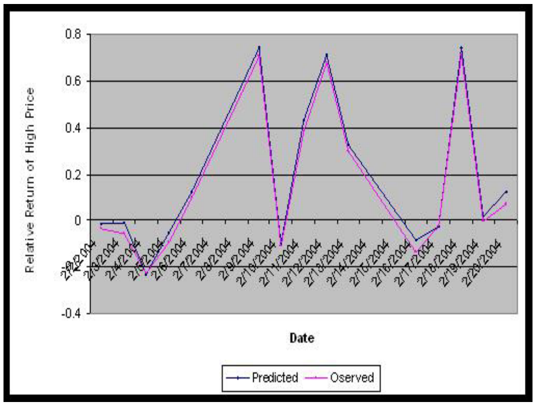

Figure 8: Predicted and actual values for the relative returns of the High price versus date

As shown in Figure 08, the direction of the predicted and the observed values are the same for many cases. Therefore, it is reasonable to state that the NN is satisfactory for predicting the direction of the High price. Furthermore, the predicted and the observed values during the considered period are very close for all except for a few days.

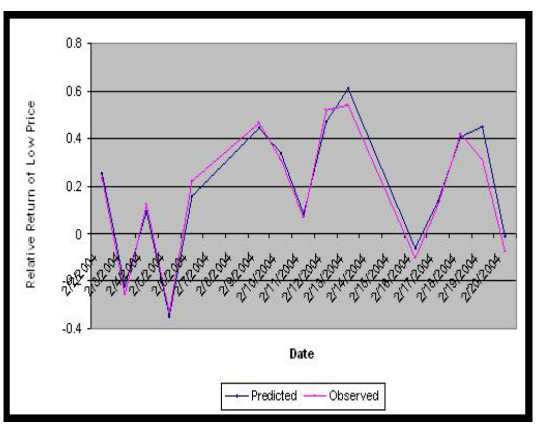

Figure 9: Predicted and actual values for the relative returns of the Low price versus date

As the above figure shows, the direction of both the predicted and actual values is the same for all days within the considered period, with the exception for a few days. Hence, it can be assume that the NN is suitable also for predicting the direction of the Low price.

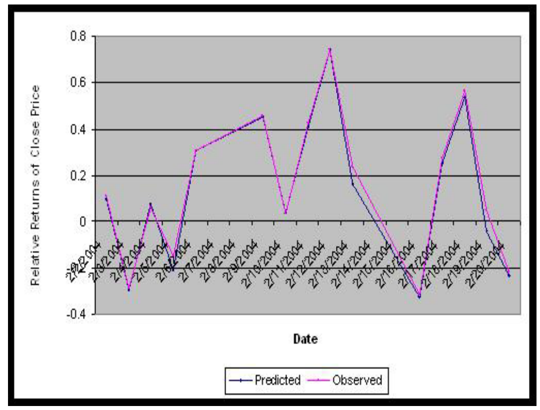

Figure 10: Predicted and actual values for the relative returns of the Close price versus date

As shown in the above figure, the NN predicts the direction of the Close price of the target market accurately, with the exception of three days. Hence, it is reasonable to assume that this network has the ability to predict the direction of the Close price accurately.

6. CONCLUSIONS

A single-layer feedforward neural network, with 33 input neurons, 3 neurons in the hidden layer, and four output neurons shows the best performance on test data. The learning rate and the momentum coefficient were 0.03 and 0.4 respectively. This network has the ability to predict the direction of the Open, High, Low, and Close prices of the AORD with accuracy of 87.29%, 80.51%, 83.05%, and 81.36%, respectively. This matter suggests that the NN approach proposed in this study is satisfactory.

Furthermore, this study revealed that the Open, High, Low, and the Close of the AORD showed approximate 6-day cycles, during the study period.

REFERENCES

Cao, L., and Tay, F. E. H. (2001). Financial forecasting using support vector machines. *Neural Computing and Applications 10*,184-192.

Chenoweth, T., and Obradovic, Z. (1995). A multi-component approach to stock market predictions, *Proceedings of the Third international conference on artificial intelligence applications on Wall Street* (pp. 74-79).

Cybenko, G. (1989). Approximation by superpositions of a sigmodial function. *Mathematics of Control, Signals and Systems, 2*, 303-314.

Egeli, B., Ozturan, M., and Badur, B. (2003). Stock Market Prediction Using Artificial Neural Networks. *Proceedings of the Hawaii International Conference on Business*, Honolulu, Hawaii, USA.

Gately, E. (1996). *Neural Networks for Financial Forecasting*. New York: John Wiley & Sons, Inc.

Gencay R., and Stengos T. (1998) Moving Average Rules, Volume and the Predictability of Security Returns with Feedforward Networks. *Journal of Forecasting 17*, 401-414

Huang, W., Nakamori, Y., and Wang, S. Y. (2005) Forecasting stock market movement direction with support vector machine. *Computational and Operational Research 32*, 1-10.

Kaastra, I., and Boyd, M. (1996). Designing a neural network for forecasting financial and economic time series. *Neurocomputing, 10*, 215-236.

Kim, S. H., and Chun, S. H. (1998). Graded forecasting using an array of bipolar predictions: Application of probabilistic neural networks to a stock market index. *International Journal of Forecasting 14*, 323-337.

Kolarik, T., and Rudorfer, G. (1994). Time Series Forecasting Using Neural Networks, *Proceedings of the International Conference on APL: the language and its application* (pp. 86-94). Antwerp, Belgium.

Leung, M. T., Daouk, H., and Chen, A. S. (2000) Forecasting stock indices: a comparison of classification and level estimation models. *International Journal of Forecasting 16*, 173-190

Malkiel, B. G. (1985). *A Random Walk Down Wall Street*. New York: W. W. Nortan & Company.

Mehrotra, K., Mohan, C. K., and Ranka, S. (1997). *Elements of Artificial Neural Networks*. Massachusetts: MIT Press.

Mizuno, H., Kosaka, M., and Yajima, H. (1998, 7 December 2001). *Application Of Neural Network To Technical Analysis Of Stock Market Prediction*. Retrieved 5th July, 2003, from the World Wide Web: www.ici.ro/ici/revista/sic98_2/art03.html

Olson, D., and Mossman, C. (2002). Neural network forecasts of Canadian stock returns using accounting ratios. *International Journal of Forecasting, 19*, 453-465.

Pan, H., Tilakaratne, C., and Yearwood, J. (2003c). Predicting the Australian Stock Market Index Using Neural Networks Exploiting Dynamical Swings and Intermarket Influences, *AI 2003: Advances in Artificial Intelligence* (pp. 327-338). Berlin, Germany: Springer-Verlag. Also in *Journal of Research and Practice in Information Technology 37*, 43-55

Qi, M. (1999). Nonlinear Predictability of Stock Returns Using Financial and Economic Variables. *Journal of Business & Economic Statistics, 17*(4), 419-429.

Qi, M., Maddala, G. S. (1999). Economic factors and the stock market: A new perspective. *Journal of Forecasting 18*, 151-166

Refenes, A. N., Zapranis, A., and Francis, G. (1994). Stock Performance Modeling Using Neural Networks: A comparative Study With Regression Models. *Neural Networks, 7*(2), 375-388.

Tilakaratne, C. D., Mammadov, M. A., and Hurst, C. P. (2006). Quantification of Intermarket Influence Based on the Global Optimization and Its Application for Stock Market Prediction. *Proceedings of the International Workshop on Integrating AI and Data Mining (AIDM'06)* http://doi.ieeecomputersociety.org/10.1109/AIDM.2006.14

Safer, A. M. (2003). A comparison of two data mining techniques to predict abnormal stock market returns. *Intelligent Data Analysis 7*, 3-13

Weiss, S. M., and Kulikowski, C. A. (1991). *Computer systems that learn*. California: Morgan Kaufmann.

White, H. (1988). Economic Prediction Using Neural networks: The Case of IBM Daily Stock Returns, *Proceedings of the Second Annual IEEE Conference on Neural Networks* (pp. 451-458). New York: IEEE Press.

Wu, Y. and Zhang, H. (1997). Forward Premiums as Unbiased Predictors of Future Currency Depreciation: A Non-parametric Analysis. *Journal of International Money and Finance 16*, 609-623

Yao, J., and Poh, H. L. (1995). Forecasting The KLSE Index Using Neural Networks, *Proceedings of the 1995 IEEE International Conference on Neural Networks* (pp. 1012-1017). Perth, Western Australia.