Bond University
ePublications@bond

Information Technology papers | Faculty of Business

1-1-2009

Designing short term trading systems with artificial neural networks

Bruce Vanstone
Bond University, Bruce_Vanstone@bond.edu.au

Gavin Finnie
Bond University, Gavin_Finnie@bond.edu.au

Tobias Hahn
Bond University, Tobias_Hahn@bond.edu.au

Follow this and additional works at: http://epublications.bond.edu.au/infotech_pubs

Part of the Computer Sciences Commons, and the Portfolio and Security Analysis Commons

Recommended Citation

Vanstone, Bruce; Finnie, Gavin; and Hahn, Tobias, "Designing short term trading systems with artificial neural networks" (2009). *Information Technology papers.* Paper 68.
http://epublications.bond.edu.au/infotech_pubs/68

This Book Chapter is brought to you by the Faculty of Business at ePublications@bond. It has been accepted for inclusion in Information Technology papers by an authorized administrator of ePublications@bond. For more information, please contact Bond University's Repository Coordinator.

Chapter #

Designing Short Term Trading Systems With Artificial Neural Networks

Bruce Vanstone[1], Gavin Finnie[2], Tobias Hahn[3]

Introduction

There is a long established history of applying Artificial Neural Networks (ANNs) to financial data sets, with the expectation of discovering financially viable trading rules. Despite the large amount of published work in this area, it is still difficult to answer the simple question, "Can ANNs be used to develop financially viable stockmarket trading systems?" Vanstone and Finnie (2007) have provided an empirical methodology which demonstrates the steps required to create ANN-based trading systems which allow us to answer this question.

In this paper, the authors demonstrate the use of this methodology to develop a financially viable, short-term trading system. When developing short-term systems, the authors typically site the neural network within an already existing non-neural trading system. This paper briefly reviews an existing medium-term long-only trading system, and then works through the Vanstone and Finnie methodology to create a short-term focused ANN which will enhance this trading strategy.

The initial trading strategy and the ANN enhanced trading strategy are comprehensively benchmarked both in-sample and out-of-sample, and the

[1] Bruce Vanstone is an Assistant Professor at Bond University, Australia (phone: +61-7-55953394; fax: +61-7-55953320; email: bvanston@bond.edu.au
[2] Gavin Finnie is a Professor at Bond University
[3] Tobias Hahn is a PhD student at Bond University

superiority of the resulting ANN enhanced system is demonstrated. To prevent excessive duplication of effort, only the key points of the methodology outlined are repeated in this paper. The overall methodology is described in detail in Vanstone and Finnie (2007), and this methodology is referred to in throughout this paper as 'the empirical methodology'.

Review of Literature

There are two primary styles of stockmarket trader, namely Systems traders, and Discretionary traders. Systems traders use clearly defined rules to enter and exit positions, and to determine the amount of capital risked. The strategies created by systems traders can be rigorously tested, and clearly understood. The alternative, discretionary trading, is usually the eventual outcome of an individual's own experiences in trading. The rules used by discretionary traders are often difficult to describe precisely, and there is usually a large degree of intuition involved. In many cases, some of the rules are contradictory – in these cases, the discretionary trader uses experience to select the appropriate rules. Despite these obvious drawbacks, however, it is commonly accepted that discretionary traders produce better financial results (2006).

For the purposes of this paper, it is appropriate to have a simple, clearly defined mathematical signal which allows us to enter or exit positions. This allows us to accurately benchmark and analyze systems.

This paper uses the GMMA as the signal generator. The GMMA is the Guppy Multiple Moving Average, as created and described by Daryl Guppy (2004), a leading Australian trader. Readers should note that Guppy does not advocate the use of the GMMA indicator in isolation (as it is used in this study), rather it is appropriate as a guide. However, the GMMA is useful for this paper, as it is possible to be implemented mechanically. In essence, any well defined signal generator could be used as the starting point for this paper.

The GMMA is defined as:

$$GMMA = \left(\begin{pmatrix} ema(3) + ema(5) \\ + ema(8) + ema(10) \\ + ema(12) + ema(15) \end{pmatrix} - \begin{pmatrix} ema(30) + ema(35) \\ + ema(40) + ema(45) \\ + ema(50) + ema(60) \end{pmatrix} \right) \tag{1}$$

Creation of the ANNs to enhance this strategy involves the selection of ANN inputs, outputs, and various architecture choices. The ANN inputs and outputs are a cut-down version of those originally described in Vanstone (2006). The original list contained 13 inputs, and this paper uses only 5. These 5 variables, discussed later in this paper, were selected as they were the most commonly discussed in the main practitioners' journal, 'The Technical Analysis of Stocks and Commodities'. Similarly, the choices of output and architecture are described in the empirical methodology paper. Again, these are only briefly dealt with here.

For each of the strategies created, an extensive in-sample and out-of-sample benchmarking process is used, which is also further described in the methodology paper.

Methodology

This study uses data for the ASX200 constituents of the Australian stockmarket. Data for this study was sourced from Norgate Investor Services (2004). For the in-sample data (start of trading 1994 to end of trading 2003), delisted stocks were included. For the out-of-sample data (start of trading 2004 to end of trading 2007) delisted stocks were not included. The ASX200 constituents were chosen primarily for the following reasons:

1. The ASX200 represents the most important component of the Australian equity market due to its high liquidity – a major issue with some previously published work is that it may tend to focus too heavily on micro-cap stocks, many of which do not have enough trading volume to allow positions to be taken, and many of which have excessive bid-ask spreads,

2. This data is representative of the data which a trader will use to develop his/her own systems in practice, and is typical of the kind of data the system will be used in for out-of-sample trading

Software tools used in this paper include Wealth-Lab Developer, and Neuro-Lab, both products of Wealth-Lab Inc (now owned by Fidelity) (2005). For the neural network part of this study, the data is divided into 2 portions: data from 1994 up to and including 2003 (in-sample) is used to predict known results for the out-of-sample period (from 2004 up to the end of 2007). In this study, only ordinary shares are considered.

The development of an ANN to enhance the selected strategy is based on simple observation of the GMMA signals. Figure #-1 shows sample buy/sell signals using the points where the GMMA signal crosses above/below zero. One of the major problems of using the GMMA in isolation is that it frequently whipsaws around the zero line, generating spurious buy/sell signals in quick succession.

One possible way of dealing with this problem is to introduce a threshold which the signal must exceed, rather than acquiring positions as the zero line is crossed. The method used in this paper, however, is to forecast which of the signals is most likely to result in a sustained price move. This approach has a major advantage over the threshold approach; namely, in a profitable position, the trader has entered earlier, and therefore, has an expectation of greater profit. By waiting for the threshold to be exceeded, the trader is late in entering the position, with subsequent decrease in profitability.

However, for the approach to work, the trader must have a good forecast of whether a position will be profitable or not. This is the ideal job for a neural network.

In figure #-2, there are a cluster of trades taken between June 2006 and September 2006, each open for a very short period of time as the GMMA whipsaws around the zero line. Eventually, the security breaks out into a sustained up trend. What is required is an ANN which can provide a good quality short-term forecast of the return potential each time the zero line is crossed, to allow the trader to discard the signals which are more likely to become whipsaws, thus concentrating capital on those which are more likely to deliver quality returns.

Figure #-2. GMMA Signals

The neural networks built in this study were designed to produce an output signal, whose strength was proportional to expected returns in the 5 day timeframe. In essence, the stronger the signal from the neural network, the greater the expectation of return. Signal strength was normalized between 0 and 100.

The ANNs contained 5 data inputs. These are the technical variables deemed as significant from the review of both academic and practitioner publications, and details of their function profiles are provided in Vanstone (2006). The formulas used to compute these variables are standard within technical analysis. The actual variables used as inputs, and their basic statistical characteristics are provided in Table #-1.

Table #-1. Technical Variables: Statistical Properties

Variable	Min	Max	Mean	StdDev
ATR(3) / ATR(15)	0.00	3.71	1.00	0.30
STOCHK(3)	0.00	100.00	54.52	36.63
STOCHK(15)	0.00	100.00	64.98	27.75
RSI(3)	0.12	100.00	58.07	25.00
RSI(15)	32.70	98.03	58.64	8.48

For completeness, the characteristics of the output target to be predicted, the 5 day return variable, are shown in Table #-2. This target is the maximum percentage change in price over the next five days, computed for every element i in the input series as:

$$\left(\frac{\left(highest(close_{i+5\ldots i+1})-close_i\right)}{close_i}\right)\times 100 \tag{2}$$

Effectively, this target allows the neural network to focus on the relationship between the input technical variables, and the expected forward price change.

Table #-2. Target Variable: Statistical Properties

Variable	Min	Max	Mean	StdDev
Target	0.00	100.00	5.02	21.83

The calculation of the return variable allows the ANN to focus on the highest amount of change that occurs in the next 5 days, which may or may not be the 5-day forward return. Perhaps a better description of the output variable is that it is measuring the maximum amount of price change that occurs within the next 5 days. No adjustment for risk is made, since traders focus on returns and use other means, such as stop orders, to control risk.

As explained in the empirical methodology, a number of hidden node architectures need to be created, and each one benchmarked against the in-sample data.

The method used to determine the hidden number of nodes is described in the empirical methodology. After the initial number of hidden nodes is determined, the first ANN is created and benchmarked. The number of hidden nodes is increased by one for each new architecture then created, until in-sample testing reveals which architecture has the most suitable in-sample metrics. A number of metrics are available for this purpose, in this study, the architectures are benchmarked using the Average Profit/Loss per Trade expressed as a percentage. This method assumes unlimited capital, takes every trade signaled, and includes transaction costs, and measures how much average profit is added by each trade over its lifetime. The empirical methodology uses the filter selectivity metric for longer-term systems, and Tharp's expectancy (1998) for shorter term systems. This paper also introduces the idea of using overall system net profit to benchmark, as this figure takes into account both the number of trades (opportunity), and the expected return of each trade on average (reward).

Results

A total of 362 securities had trading data during the test period (the ASX200 including delisted stocks), from which 11,897 input rows were used for training. These were selected by sampling the available datasets, and selecting every 25th row as an input row.

Table #-3 reports the Overall Net System Profit, Average Profit/Loss per Trade (as a percentage), and Holding Period (days) for the buy-and-hold naïve approach (1st row), the initial GMMA method (2nd row), and each of the in-sample ANN architectures created (subsequent rows). These figures include transaction costs of $20 each way and 0.5% slippage, and orders are implemented as day+1 market orders. There are no stops implemented in in-sample testing, as the objective is not to produce a trading system (yet), but to measure the quality of the ANN produced. Later, when an architecture has been selected, stops can be determined using ATR or Sweeney's(1996) MAE technique.

The most important parameter to be chosen for in-sample testing is the signal threshold, that is, what level of forecast strength is enough to encourage the trader to open a position. This is a figure which needs to be chosen with respect to the individuals own risk appetite, and trading requirements. A low threshold will generate many signals, whilst a higher threshold will generate fewer. Setting the threshold too high will mean that trades will be signalled only rarely, too low and the trader's capital will be quickly invested, removing the opportunity to take higher forecast positions as and when they occur.

For this benchmarking, an in-sample threshold of 5 is used. This figure is chosen by visual inspection of the in-sample graph in Figure #-3, which shows a breakdown of the output values of a neural network architecture (scaled from 0 to 100) versus the average percentage returns for each network output value. The percentage returns are related to the number of days that the security is held, and these are shown as the lines on the graph. Put simply, this graph visualizes the returns expected from each output value of the network and shows how these returns per output value vary with respect to the holding period. At the forecast value of 5 (circled), the return expectation rises above zero, so this value is chosen.

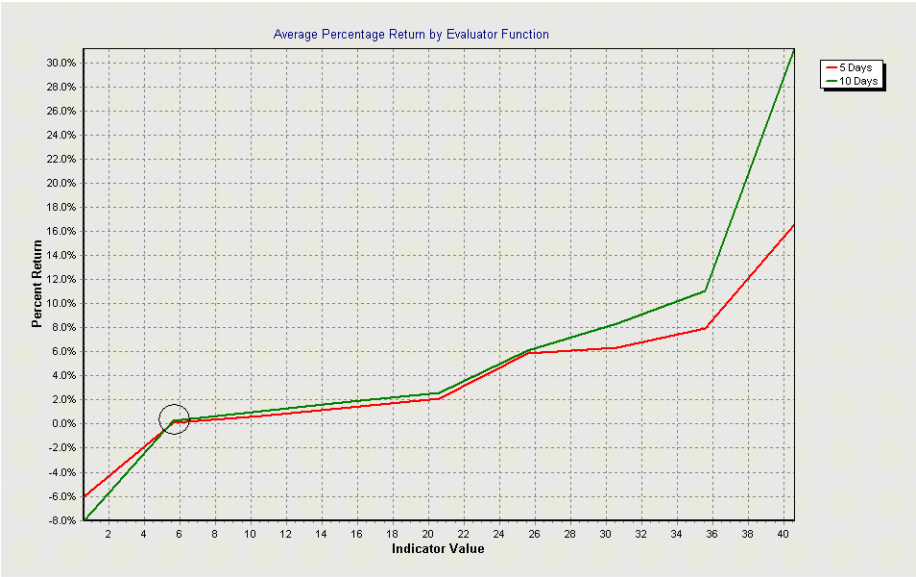

Figure #-3. In-sample ANN function profile

Table #-3. In-Sample characteristics

Strategy (In-Sample Data)	Overall Net System Profit	Profit/Loss per Trade (%)	Holding Period (days)
Buy-and-hold naïve approach	1,722,869.39	94.81	2,096.03
GMMA alone	632,441.43	1.09	35.30
ANN – 3 hidden nodes + GMMA	878,221.68	2.32	46.15
ANN – 4 hidden nodes + GMMA	1,117,520.33	3.69	59.20
ANN – 5 hidden nodes + GMMA	353,223.61	3.00	42.64

As described in the empirical methodology, it is necessary to choose which ANN is the 'best', and this ANN will be taken forward to out-of-sample testing. It is for this reason that the trader must choose the in-sample

benchmarking metrics with care. If the ANN is properly trained, then it should continue to exhibit similar qualities out-of-sample which it already displays in-sample.

From the above table, it is clear that ANN – 4 hidden nodes should be selected. It displays a number of desirable characteristics – it shows the highest level of Profit/Loss per Trade. Note that this will not necessarily make it the best ANN for a trading system. Extracting good profits in a short time period is only a desirable trait if there are enough opportunities being presented to ensure the traders capital is working efficiently.

Therefore, it is also important to review the number of opportunities signalled over the 10-year in-sample period. This information is shown in Table #-4.

Table #-4. Number of Trades Signalled

Strategy (In-Sample Data)	Number of trades signaled
Buy-and-hold naïve approach	362
GMMA alone	11,690
ANN – 3 hidden nodes + GMMA	7,516
ANN – 4 hidden nodes + GMMA	6,020
ANN – 5 hidden nodes + GMMA	2,312

Here the trader must decide whether the number of trades signalled meets the required trading frequency. In this case, there are likely to be enough trades to keep an end-of-day trader fully invested.

This testing so far covered data already seen by the ANN, and is a valid indication of how the ANN should be expected to perform in the future. In effect, the in-sample metrics provide a framework of the trading model this ANN should produce.

Table #-5 shows the effect of testing on the out-of-sample ASX200 data, which covers the period from the start of trading in 2004 to the end of trading in 2007. These figures also include transaction costs and slippage, and orders are implemented as next day market orders.

This was a particularly strong bull market period in the ASX200.

Table #-5. Out-of-Sample Performance

Strategy (Out-of-Sample Data)	Overall Net System Profit	Profit/Loss per Trade (%)
GMMA alone	622,630.01	3.88
ANN – 4 hidden nodes + GMMA	707,730.57	10.94

Although there appears a significant difference between the GMMA, and the ANN enhanced GMMA, it is important to quantify the differences statistically. The appropriate test to compare two distributions of this type is the ANOVA test (see supporting work in Vanstone (2006)). The results for the ANOVA test are shown in Table #-6 below.

Table #-6. ANOVA Comparison

ANOVA	GMMA	GMMA + 4 hidden nodes
Number of observations	3,175	1,283
Mean	196.10	551.62
Std Dev	1,496.99	2,483.64
95% Confidence Internal of the mean – lower bound	144.01	415.59
95% Confidence Internal of the mean – upper bound	248.19	687.65

The figures above equate to an F-statistic of 34.26, (specifically, $F(1,4456) = 34.261$, $p=0.00$ ($p<0.05$)), which gives an extremely high level of significant difference between the two systems.

Conclusions

The ANN out-of-sample performance is suitably close to the ANN in-sample performance, leading to the conclusion that the ANN is not curve fit, that is, it should continue to perform well into the future. The level of significance reported by the ANOVA test leads to the conclusion that the ANN filter is making a statistically significant improvement to the quality of the initial GMMA signals.

The trader now needs to make a decision as to whether this ANN should be implemented in real-life.

One of the main reasons for starting with an existing successful trading strategy is that it makes this decision much easier. If the trader is already using the signals from a system, and the ANN is used to filter these signals, then the trader is still only taking trades that would have been taken by the original system. The only difference in using the ANN enhanced system is that trades with low expected profitability should be skipped.

Often in trading, it is the psychological and behavioural issues which undermine a traders success. By training ANNs to support existing systems, the trader can have additional confidence in the expected performance of the ANN.

Finally, Figure #-4 shows the same security as Figure #-2. The ANN has clearly met its purpose of reducing whipsaws considerably, which has resulted in the significant performance improvement shown in Table #-3 and Table #-5.

Of course the result will not always be that all whipsaws are removed. Rather, only whipsaws which are predictable using the ANN inputs will be removed.

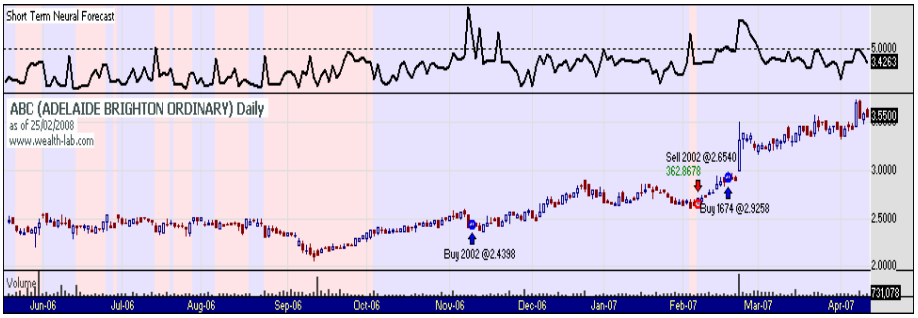

Figure #-4. GMMA signals filtered with an ANN

References

(2004). "Norgate Premium Data." Retrieved 01-01-2004, from www.premiumdata.net.

(2005). "Wealth-Lab." from www.wealth-lab.com.

Elder, A. (2006). Entries & Exits : visits to sixteen trading rooms. Hoboken: NJ, John Wiley and Sons.

Guppy, D. (2004). Trend Trading. Milton, QLD, Wrightbooks.

guppytraders.com. "Guppy Multiple Moving Average." Retrieved 04-05-2007, from www.guppytraders.com/gup329.shtml.

Sweeney, J. (1996). Maximum Adverse Excursion: analyzing price fluctuations for trading management. New York, J. Wiley.

Tharp, V. K. (1998). Trade your way to Financial Freedom. NY, McGraw-Hill.

Vanstone, B. (2006). Trading in the Australian stockmarket using artificial neural networks, Bond University. PhD.

Vanstone, B. and G. Finnie (2007). "An Empirical Methodology for developing Stockmarket Trading Systems using Artificial Neural Networks." Expert Systems with Applications In-Press (DOI: http://dx.doi.org/10.1016/j.eswa.2008.08.019).