2014 International Joint Conference on Neural Networks (IJCNN)
July 6-11, 2014, Beijing, China

Beating The S&P 500 Index - A Successful Neural Network Approach

Dr. Mininder Sethi
Centre for Financial Computing
University College London
London, United Kingdom
m.sethi@cs.ucl.ac.uk

Prof. Philip Treleaven
Centre for Financial Computing
University College London
London, United Kingdom
p.treleaven@ucl.ac.uk

Dr. Sebastian Del Bano Rollin
Centre for Financial Computing
University College London
London, United Kingdom
s.delbanorollin@cs.ucl.ac.uk

***Abstract*— The systematic trading of equities forms the basis of the asset management industry. Analysts are trying to outperform a passive investment in an index such as the S&P 500 Index. However, statistics have shown that most analysts fail to consistently beat the index. A number of Neural Network based methods for detecting trading opportunities on Futures contracts on the S&P 500 Index have been published in the literature. However, such methods have generally been unable to demonstrate sustained performance over a significant period of time. The authors of this paper show, through the application of over ten years of experience in quantitative modelling and trading, a different type of Neural Network approach to beating the S&P 500 Index. Rather than trading Futures contracts, it is shown that by using Neural Networks to intelligently select just a handful of stocks a performance significantly in excess of a buy and hold position on the S&P 500 Index could have been achieved over a seven year period. The effect of transaction costs is also considered.**

Keywords*— Advanced Computational Intelligence for Algorithmic Trading, Applications of Neural Networks for Financial Modelling and Forecasting.

I. INTRODUCTION

The use of Neural Network based methods to find trading opportunities on the S&P 500 Index or on Futures contracts on the Index has been considered in the literature [1-5]. The aim is typically to outperform a benchmark buy and hold position on the S&P 500 Index, thus achieving something that over 65% of large-cap fund managers could not [6]. Such Neural Network based approaches have shown varying degrees of success, but they have generally been unable to show sustained profitability over a significant period of time. These methods have typically used as input features Technical Analysis Indicators [7] often without an economic rationale. A method that incorporates real trading experience and is based on a set of input features, with a tractable economic basis, would be preferable.

In a previous publication [8] the authors of this paper have presented a new Neural Network framework that could be used to detect both long and short trading opportunities at the single stock level. The framework was based upon a justified economic rational. This paper is a companion paper to [8] and uses the same structure to allow easy comparison. It shows that the techniques could be applied to form a long only portfolio of the constituent stocks of the S&P 500 Index and that such a portfolio would have significantly outperformed the index.

The premise of the method is that individual stocks should generally be considered as efficiently priced. However occasional anomalies will arise and these will give way to identifiable short term trends. For any particular stock the method will for the most time determine that no identifiable trend exists. However, for a wide pool of stocks, such as the constituents of the S&P 500 Index, the Neural Network based framework will find sufficient long only trading opportunities to show that it would have significantly beaten a buy and hold position in the S&P 500 Index over an extensive seven year time period even with the inclusion of transaction costs.

II. TRADING SIGNAL GENERATION

In this section a method of stock closing price data pre-processing is presented. This section is important as it provides an economic justification of why the chosen input features are used with the Neural Network.

A compact representation of the price trend over some time period can be achieved through just two metrics, a Short Term Efficiency Level and an Average Efficiency Level. The Efficiency Level measure is similar to that from Kaufman [9]. Assume that market data has been regularly sampled with the stock closing price at the nth time sample for a stock with ticker symbol TCK being represented as $S_{TCK}[n]$. The Short Term Efficiency Level at timestamp n can then be defined

$$\gamma_{TCK}[n] = \frac{|S_{TCK}[n] - S_{TCK}[n-K]|}{\sum_{k=0}^{K-1} |S_{TCK}[n] - S_{TCK}[n-k-1]|} \tag{1}$$

Where |.| is the absolute value operator. In summary the Short Term Efficiency Level calculation looks back from timestamp n over a window of K observations and provides the ratio of the absolute end to end stock price change to the sum of the absolute day to day changes over that window. The Short Term Efficiency Level can be seen as a ratio of Signal to Signal Plus Noise and hence has a tractable meaning in an Engineering sense. In an economic sense the Short Term Efficiency Level can be seen as a measure of trend, a level close to 1 symbolizing a straight line movement and hence a strong trend and a level close to zero symbolizing day to day movement within a period of K observations but with no end to end movement, hence no trend. The value for K could be determined through a back testing process and K could be made stock dependent and adaptive. For simplicity the value $K = 10$ is initially taken for all stocks.

As well as having interest in the Short Term Efficiency Level (trend level), the evolution of the trend would be expected to contain information. Such an evolution could be observed through an Efficiency Vector which could be defined

$$\vec{\gamma}_{TCK}[n] = [\gamma_{TCK}[n], \gamma_{TCK}[n-1], \dots, \gamma_{TCK}[n-L]] \qquad (2)$$

Such a vector would keep track of the Short Term Efficiency Levels over the current and L preceding timestamps. In the interest of model compactness it is instead proposed to consider only an Average Efficiency Level which is defined

$$\bar{\gamma}_{TCK}[n] = \frac{1}{L}\sum_{l=0}^{L-1} \gamma_{TCK}[n-l] \qquad (3)$$

The Average Efficiency Level is used as a proxy for the average trend in the stock price over the preceding period of L timestamps. The value for L could be determined through a back testing process and L could be made stock dependent and adaptive. For simplicity the value $L = 20$ is initially used.

A Neural Network based system that considers just two input features, the Short Term Efficiency Level and the Average Efficiency Level, may appear simplistic. However, the system is built upon a sound economic premise and as such represents an improvement over methods that utilise many Technical Analysis Indicators without a rational basis.

In order to generate a training set the Profit and Loss (P&L) of a trade entered to follow the Short Term Trend Direction of a stock can be considered. The Short Term Trend Direction at timestamp n for the stock with ticker TCK can be defined as

$$D_{TCK}[n] = \operatorname{sgn}(S_{TCK}[n] - S_{TCK}[n-K]) \qquad (4)$$

Where sgn (.) is the sign operator and takes the value of $+1$ if its operand is greater than or equal to zero and takes the value of -1 otherwise. A value of $D_{TCK}[n] = 1$ then corresponds to the case that the stock is seen to be in a short term uptrend and the value of $D_{TCK}[n] = -1$ corresponds to the case that the stock is seen to be in a short term downtrend. The indicator $D_{TCK}[n]$ provides only directional information, the Short Term Efficiency Level $\gamma_{TCK}[n]$ provides a measure of the recent strength of the trend.

The training set is generated such that the nth training sample is $\{\{\gamma_{TCK}[n], \bar{\gamma}_{TCK}[n]\}, P_{TCK}[n]\}$ where $\gamma_{TCK}[n]$ and $\bar{\gamma}_{TCK}[n]$ are the Short Term and Average Efficiencies and $P_{TCK}[n]$ is the Categorized Trade P&L for placing a trade in the direction $D_{TCK}[n]$. The Categorized Trade P&L, $P_{TCK}[n]$, will be the eventual Neural Network output and will be able to take one of 9 possible discrete category values in the range $[-1,1]$ as in Table I. The Trade Categorisation Threshold ω in Table I could be made stock dependent and adaptive, for simplicity a value of $\omega = 0.30\%$ is taken for all stocks.

The Categorized Trade P&L corresponds to the trading profit from a trade that follows the Short Term Trend Direction $D_{TCK}[n]$. In the case that $P_{TCK}[n] \leq -0.25$ it is implied that a positive profit in excess of $\omega = 0.30\%$ could have been made by trading in the opposite direction to $D_{TCK}[n]$. For the purpose of generating the nth training sample the value of $P_{TCK}[n]$ is generated by considering a trade of direction $D_{TCK}[n]$ entered at timestamp n, where the trade is exited after $j \geq 1$ days when the earliest of three exit criteria is satisfied (i) $\gamma_{TCK}[n+j] < \gamma_{TCK}[n]$, (ii) $\bar{\gamma}_{TCK}[n+j] < \bar{\gamma}_{TCK}[n]$ or $j = 3$. The first two criteria correspond to a weakening of the Short Term or Average Efficiency Levels below their levels at trade initiation, the final criteria imposes a maximum holding time of 3 days. The first two trade exit criteria are naturally intuitive and the third imposes a pseudo risk limit that prevents any trade being held for too long.

The motivation for the categorization illustrated in Table I is to establish an expected P&L range under which no trades would be placed (category zero) and to establish a set of expected positive and negative P&L ranges which can be used to determine if the Short Term Trend Direction should be followed or countered. Having a range of values as opposed to a categorization such as $\{-1,0,1\}$ allows for the classification of high P&L outliers (categories -1 and 1) and also allows for the possible incorporation of leverage in the case that the expected P&L magnitude is large.

III. APPLICATION OF A NEURAL NETWORK

Given that there are only two input features a 20 Neuron Neural Network is found to be sufficient. Since the input feature space is limited to just two dimensions, a

Table I. Trade P&L Categories

Category	P&L Range
-1.00	$-\infty < P\&L \leq -7\omega$
-0.75	$-7\omega < P\&L \leq -5\omega$
-0.50	$-5\omega < P\&L \leq -3\omega$
-0.25	$-3\omega < P\&L \leq -\omega$
0.00	$-\omega < P\&L \leq \omega$
0.25	$\omega < P\&L \leq 3\omega$
0.50	$3\omega < P\&L \leq 5\omega$
0.75	$5\omega < P\&L \leq 7\omega$
1.00	$7\omega < P\&L \leq \infty$

straightforward visualization of the training set is possible. Figure 1 shows a training set of 750 training samples for McDonalds (ticker MCD) generated on the daily closes between May 2010 and April 2013. The training set is noisy as expected. In Figure 2 the outputs generated by a Neural Network trained with this data are shown across discretely sampled intervals of the range of possible values of the two input feature space $\{\gamma_{MCD}[n], \bar{\gamma}_{MCD}[n]\}$. From Figure 2 it can be seen that in regions of the feature space where little training data was available the Neural Network has made inferences of an expected large and negative P&L for a trade that would follow the Short Term Trend Direction $D_{MCD}[n]$. It is unsurprising that the output would be somewhat random in such regions given the limited span of the training data.

It may be argued that the issue shown in Figure 2 can be ignored. If a training set is devoid of samples that fall in some region of feature space, it may then be expected in practice to confront such regions of feature space with a low probability and therefore the Neural Network inferences are not so important. However, a form of heuristic regularization that can deal with such an issue is to append to the training set a subset of biasing training samples that bias the Neural Network output towards the category that makes the decision for no trade (category zero). An example of such a modified training set is shown in Figure 3 where zero output training samples have been regularly placed in the feature space. The cleaner Zero Biased Neural Network output is shown in Figure 4. In practice the procedure of zero biasing a training set will lead the Neural Network away from making a decision to trade. However, given a limited amount of investable money a missed opportunity for any particular stock would allow trading on other stocks. A trading decision at some post training timestamp m can be made by evaluating the Neural Network output given the prevailing Efficiency Levels. The Neural Network output can be converted to a P&L category as

$$P_{TCK}[m] = max\left(-1, min\big(1, 0.25 * round(4 * Z_{TCK}[m])\big)\right) \qquad (5)$$

where $Z_{TCK}[m]$ is the output for inputs $\{\gamma_{TCK}[m], \bar{\gamma}_{TCK}[m]\}$ and round(.) is the round to nearest integer operator.

IV. FRAMEWORK IN PRACTICE

In this section results are presented to demonstrate the effectiveness of the proposed Neural Network framework. The generation of trading signals for individual single stocks is

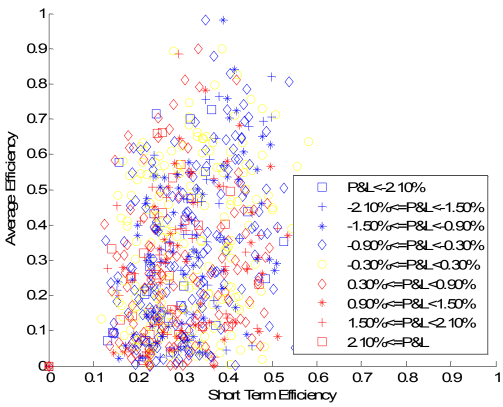

Fig. 1. Example Training Set

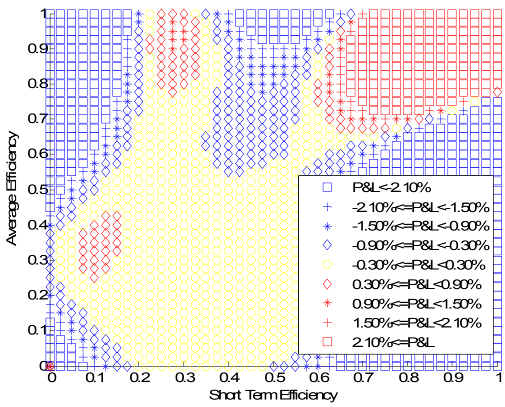

Fig. 2. Trained Neural Network Output

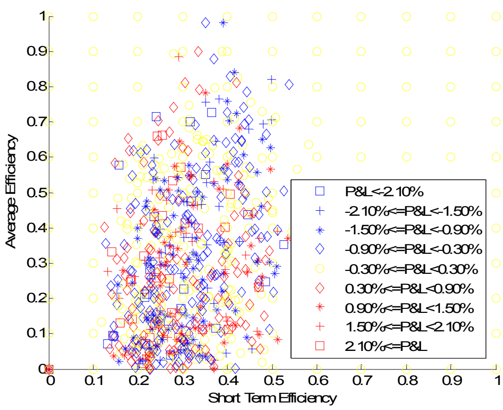

Fig. 3. Zero Appended Example Training Set

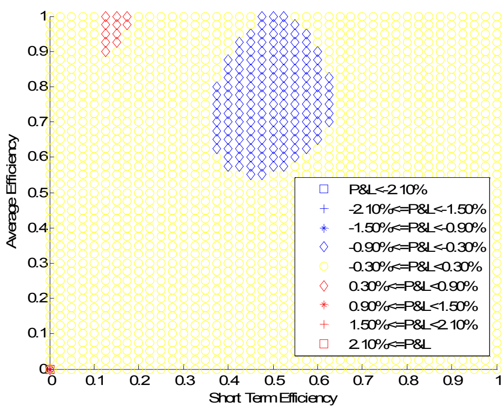

Fig. 4. Retrained Neural Network Output

considered first. This is followed by the creation of a Long Only Trading Portfolio that is shown to outperform the benchmark S&P 500 Index, even with the inclusion of transaction costs. A universe of 100 stocks which were the highest weighted constituents of the S&P 500 Index is used. An extensive time period of around 7 years spanning from May 2006 to April 2013 is considered. Such a period encompasses a range of conditions including the 2008 Global Financial Crisis [10], for reference the VIX 'investor fear gauge' [11] over the period was in the range from 10.4 to 59.9, illustrating a wide market volatility range. Closing price data has been sourced from Bloomberg and has been back adjusted for stock splits and dividends. For each stock the Neural Network is retrained every $M = 250$ business days using data for the preceding $N = 750$ business days. Average test results for long only trades generated across the universe of 100 stocks are shown in Table II. From Table II it can be seen that on average there are just over 92 trades placed per stock from a possible 1750 trades. The average trade duration is around 1.50 days and therefore on average, for each stock, there is no investment being made around 92% of the time. The framework looks only to trade when conditions are suitable. There is a high average positive profit per trade.

A Long Only Portfolio can be created on any trading day for which at least one Buy trade can be identified. Where multiple buying opportunities can be identified the available trading proceeds are shared equally amongst the trading opportunities. This method of portfolio construction is simplistic and provides a lower bound to the potential performance of the method. Improved performance would be expected if advanced portfolio construction techniques were applied. The performance of a Long Only Portfolio based upon trades generated by the proposed Neural Network Framework is shown in Figure 5. In the same figure the performance of a Buy and Hold position in the benchmark S&P 500 Index is also shown alongside the performance of a Buy and Hold position in an equally weighted basket of the same 100 stocks that are used with the proposed method. For the proposed method a transaction cost of 0.02% (2 basis points) is charged per trade. The proposed Long Only Portfolio has shown an Average (geometric) Annualized Return of 16.8% over the seven year test period in which S&P 500 Index has been relatively flat on an end to end basis.

Table II. Average Test Results

Measure	Average Over 100 Stocks
Number of Trades	92.2
Number of Profitable Trades	52.6
% Profitable Trades	57.0%
Total P&L Over 7 Years	+ 29.7%
Profit Per Trade	+ 0.32%
Total Days in Trade	137.9
Number of Days Per Trade	1.50

V. CONCLUSIONS AND FURTHER WORK

In this paper a different type of Neural Network approach for the outperformance of the S&P 500 Index has been presented. It has been shown that by intelligently selecting just a handful of the constituent stocks for which a short term trend can be identified a dynamic portfolio that would have significantly outperformed a Buy and Hold position in the S&P 500 Index could have been constructed. The framework offers room for improvement. A number of system parameters have been set manually. The equally weighted portfolio approach is also naïve and in a future piece of work a novel method for intelligent asset allocation which leads to improved performance will be presented alongside results across a range of Global stock markets.

REFERENCES

[1] R.R. Trippi and D. DeSieno. "Trading Equity Index Futures With a Neural Network", The Journal of Portfolio Management, 1992, pp. 27-33.
[2] T. Chenoweth, Z. Obradovic and S.L. Sauchi "Embedding Technical Analysis Into Neural Network Based Trading Systems". Applied Artificial Intelligence. 1996, pp. 523-542.
[3] Y. Zhang and L. Wu. "Stock Market Prediction of S&P 500 via Combination of Improved BCO Approach and BP Neural Network". Expert Systems with Applications, 2009, pp. 8849-8854.
[4] D Wood and S Vasilyev. "A Projection of S&P 500 Index Using Artificial Neural Network." European Journal of Industrial and System Engineering, 2012, pp. 1-7.
[5] S. Taghi; A. Niaki and S. Hoseinzade. "Forecasting S&P 500 Index Using Artificial Neural Networks and Design of Experiments", Springer Journal of Industrial Engineering International, 2013, pp. 1-9.
[6] http://www.forbes.com/sites/rickferri/2012/10/11/indexes-beat-active-funds-again-in-sp-study/
[7] C.D. Kirkpatrick and J.R. Dahlquist, "Technical Analysis: The Complete Resource for Financial Market Technicians". Prentice Hall, 2010.
[8] M. Sethi, P. Treleaven and S. Del Bano Rollin, " A New Neural Network Framework for Profitable Long-Short Equity Trading". In Press. Accepted to The 2014 International Conference on Computational Science and Computational Intelligence (CSCI). March 2014.
[9] P.J. Kaufman, "Alpha Trading: Profitable Strategies That Remove Directional Risk". Wiley, 2011.
[10] J. Crotty. "Structural Causes of the Global Financial Crisis: A Critical Assessment of the New Financial Architecture", Cambridge Journal of Economics, 2009, pp. 563-580.
[11] R.E. Whaley. "The Investor Fear Gauge". The Journal of Portfolio Management, 2000, pp. 12-17.

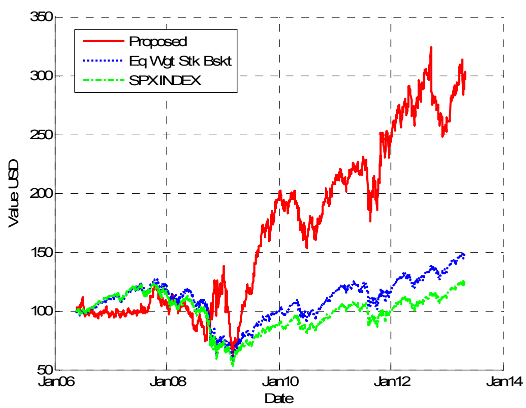

Fig. 5. Performance Against Benchmarks