

Evolving short-term trading strategies using Genetic Programming

Nils Svängård and Peter Nordin
Chalmers University of Technology
Complex Systems Group
SE-41296 Göteborg
mp00nisv@dd.chalmers.se; nordin@fy.chalmers.se

Stefan Lloyd and Clas Wihlborg
The Gothenburg School of Economics & Commercial Law
Department of Economics
Box 640; SE-405 30 Göteborg
stefan.lloyd@hgus.gu.se; cw.fi@cbs.dk

Abstract - We have used a linear Genetic Programming system with a multitude of different quotes on financial securities as input in order to evolve an intraday trading strategy for an individual stock, attempting to outperform a simple buy and hold strategy over the same period of time.

I. INTRODUCTION

During the 1990'ies equity markets experienced an enormous increase in capital inflow. The number of retail investors who engaged in day trading and individual stock picking soared. In light of these events a wide range of Internet-based brokers have emerged with an even wider range of analytical tools with which to analyze the stock market. With this ever-increasing flow of information it has become very hard for institutional, as well as individual investors to find relevant information use it efficiently, despite the multitude of analytical tools available today.

Parallel with the evolution of the financial markets there has been significant developments in the field of Artificial Intelligence, especially in soft computing using Neural Networks and Evolutionary Algorithms. We use a flavor of Genetic Programming that was initiated in 1992 by John Koza [1].

This paper intends to bring together the problem of discerning and analyzing financial data and the recent developments in the field of genetic programming. We use a GP engine to analyze financial time series with a granularity of one minute for the purpose of finding a successful trading strategy for a single stock.

II. BACKGROUND

An issue that is the subject of intense debate among academics and financial professionals is the Efficient Market Hypothesis (EMH) [2]. The EMH states that at any given time, security prices fully reflect all available information. Most individuals that buy and sell securities (stocks in particular), do so under the assumption that the securities they are buying are worth more than the price that they are paying, while on the other hand the securities that they are selling are worth less than the selling price. But if markets are efficient and current prices fully reflect all information, then buying

and selling securities in an attempt to outperform the market will effectively be a game of chance rather than skill.

The debate about efficient markets has resulted in several empirical studies attempting to determine whether specific markets are in fact 'efficient', and if so to what degree [3]-[4]. Early tests of the EMH focused on technical analysis and it is technical analysts whose very existence seems most challenged by the EMH.

In fact, the vast majority of studies of technical theories have found the strategies to be of little use in predicting securities prices [5].

Nevertheless, numerous stock market anomalies have been uncovered that seem to contradict the efficient market hypothesis. The search for anomalies is effectively the search for systems or patterns that can be used to outperform passive buy-and-hold strategies.

Despite the vast amount of research rejecting technical analysis as a trading strategy there are still market practitioners who use it extensively and who also claim to be successful [6]-[7].

There have also been studies that challenge the hypothesis that future price movements cannot be predicted from past data [8]-[10].

Though many empirical tests have rejected technical analysis using the established methods, recent research in adaptive systems, and Genetic Programming in particular, have proved to be harder to reject. Tests using GP to evolve trading strategies as a form of technical analysis have proven successful [11]-[14].

Furthermore, previous research has favored daily prices, which not only requires a very long observation period but also disregard the information that lies in intra day stock quotes. Iba & Sasaki used per minute quotes, from the Nikkei 225 between the 1st of April 1993 and the 30th of September 1993 and were successful in predicting price movements for the Nikkei 225 index using Genetic Programming [15].

The purpose of this paper is to use a Genetic Programming system with a multitude of different quotes on financial securities as input in order to evolve a trading strategy for an individual stock (Nokia) that outperform a simple buy and hold strategy, over the same period of time.

The purpose of this paper is not to challenge the Efficient Market Hypothesis (EMH), nor does it lie within the scope of it to compare the efficiency of genetic programming to any other method of time series analyses.

III. METHOD

The Genetic Programming system uses linear genomes and offer a wide range of different parameters that control the working conditions. See Table 1 for the basic parameter setup used in this paper.

The functional set contains the basic arithmetic and boolean operators as well as conditional statements, trigonometric, logarithmic and exponential functions.

TABLE 1
PARAMETER SETUP

Parameter	Value
Population size	10000
Mutation frequency	95%
Crossover frequency	95%
Freq. of Homologous Crossover	90%

A. Data


The terminal set consisted of a multitude of financial instruments collected between the 5th of April and the 9th of May in 2001. Every minute the instruments were sampled from a Reuters First terminal that provided real-time information from different exchanges. See Table 2 for a complete list of the financial instruments.

TABLE 2
FINANCIAL INSTRUMENTS

LME b Last ADR (Nasdaq)	EUR 10 Y Notes
LME b Bid ADR (Nasdaq)	General Index
LME b Ask ADR (Nasdaq)	OMX Index
Nokia Last ADR (NYSE)	FTSE 100 Index
Nokia Bid ADR (NYSE)	DAX 100 Index
Nokia Ask ADR (NYSE)	LME b Bid Size
Nasdaq 100	LME b Bid
GBP 1 Year deposit	LME b Ask
Stibor 1 Month	LME b Ask size
Stibor 1 Year	LME b Time
EUR spot (FX)	Nokia Bid size
CHF spot (FX)	Nokia Bid
GBP spot (FX)	Nokia Ask
JPY spot (FX)	Nokia Ask size
USD spot (FX)	Nokia Time
US 30 Y Govt Bond	LME b Last
Swedish 10 Y Govt Bond	Nokia Last
GB 10 Y Govt Bond	DJIA
MSCI European Consumer Discretionary Index	MSCI European Consumer Staples Index
MSCI European Health Care	MSCI European Utilities
MSCI European Energy Index	SEK 1 Year deposit
MSCI European Finance Index	EUR spot deposit
MSCI European Industry Index	EUR 1Year deposit
MSCI European IT Index	USD spot deposit
MSCI European Materials Index	GBP spot deposit
MSCI European Telecom Index	SEK spot deposit
1 Year USD Index deposit	

In Figure 1, the terminal set and the out of sample data is displayed in a graph based on the sell price of Nokia. The graph is split into 3 fairly equal slices, where the first third is the training data (approximately 7 trading days), the middle part the validation data (approximately 7 trading days) and the last third is the out of sample data (approximately 6 trading days).

Figure 1
Timespan of input data



The total number of observation points was 8045 and the out of sample validation of the trading agent consisted of 3847 observation points between the 10th and the 18th of May.

In addition we introduced six more inputs, consisting of the moving average (MA) over 5, 10 and 15 minutes. The MA variables are based on the sell price of both Nokia and Ericsson.

B. Preprocessing of Data

In previous research there has been a certain degree of preprocessing of the terminal set. Iba & Sasaki used per minute sampled quotes of the Nikkei 225 index, all of which they normalized to a value between 0 and 1 before submitting them to their GP system [15]. The degree of pre-processing depends on what type of GP system is available and on what the overall objective of the project is.

Considering the fact that the GP engine cannot see the changes in the variables as time progresses, we concluded that a certain degree of pre-processing is imperative for the success of our project. By converting the terminal set into relative measures for all transaction data, using equation 1, the GP system can learn that a positive observation in $t+1$ indicates an increase in the variable's value and vice versa.

$$v_t / v_{t-1} - 1 \quad (1)$$

The only variable that is not converted into a relative measure is the time observation. All other data is made into relative terms except the spread of both Ericsson and Nokia, which appear both as a relative and an absolute measure.

C. Feedback Variables

The ability to transport information over time is a prerequisite when searching for patterns in time series. In order to facilitate the ability to analyze lagged data, we constructed two feedback variables:

- 1) The first feedback variable carries forward the previous output, thus making the agent aware of its action one minute ago.
- 2) The second variable shows the relative performance of the share price since the last transaction was made. This gives the system a tool with which to analyze the current gain or loss situation, and whether it should capitalize now or in the next period.

D. The Fitness Criteria

We concluded that the easiest and most intuitive way of defining *fitness* must be the agent's ability to earn money. This also simplifies comparing the active trading strategy with the simple buy/hold strategy, which we have as benchmark.

The agent will, as a portfolio manager, use the initial cash installment to trade shares according to its evolved strategy. The aim for each agent is to maximize the value of the portfolio under the following conditions:

- 1) The agent can reinvest the proceeds from successful transactions but only blocks of shares can be purchased. If the cash balance does not equal an even number of blocks, the system will round off downwards and leave the excess cash in the account.
- 2) For simplicity in our calculations, no accrued interest is considered for the inactive cash in the account even though that would be the fact in real life. As the time frame for measuring the trading agent's performance is only a few days, the effective interest earned, during those days, will be negligible and not affect the general results.
- 3) The agent has to buy at the ask price and sell at the bid price.

The actual fitness value of an agent is defined as

$$v_0 / v_n \quad (2)$$

where:

v_n is the value of the portfolio at the end of the data set, and v_0 is the initial value of the portfolio.

Thus a fitness below 1.0 indicates that the agent is making money and a fitness higher than 1.0 that money is lost.

E. The Market Model

At the outset we defined three different states that the system will use to decide which action to take. Each state is represented by the return-value (R) of the agent. See Table 3 for a description of the states, expected market trend, and the portfolio action.

TABLE 3
RETURN VALUES AND MARKET TRENDS

State	Market Trend	Action
R < 0	Downward	Sell shares
R = 0	None	None
R > 0	Upward	Buy shares

An agent will at each point in time analyze the present data and distinguish if there is a trend in the market and return a R-value that reflects the appropriate action to embark upon.

The fitness function cannot distinguish how many shares to buy or sell on a given action command. Instead the agent will purchase the maximum amount of shares allowed, given the amount of cash in the account. Conversely a signal to sell will trigger the agent to divest all of its holdings.

In order to catch the effects of transaction costs, we chose to replicate the costs of having an account with a broker. However, tax on gross proceeds from selling a financial security will not be considered. Neither does the account in question allow the trading agent's to either short-sell a stock or go on a margin.

See Table 4 for the costs associated with transactions in the system.

TABLE 4
TRANSACTION COSTS


Relative Commission	0.1%
Minimum Commission	SEK 150
Maximum Commission	SEK 500

IV. RESULTS

The actual trading strategy will not be disclosed as it is several hundreds of instructions long. At the readers request we will provide it.

The fitness values for both training and validation can be viewed in the output chart in Figure 2. The chart shows how training and validation fitness converge relatively well over the succession of tournaments, i.e. when training fitness improves so does the validation fitness. This means that an improvement in the training fitness is of a general character as it applies to data not seen during the training procedure.

Figure 2
Training and Validation Fitness



In order to compare the trading agent and the passive buy/hold strategy we calculated the risk-adjusted return using equation 3. Where R_p is the relative return of the portfolio and σ_p is the standard deviation of the same portfolio. This ratio measures the amount of return per unit of risk, so the strategy with the highest ratio should be the best choice. The concept of risk is a complicated issue and we do not argue that the risk-adjusted return incorporates all the relevant risk in trading securities.

$$R_p / \sigma_p \quad (3)$$

The aggregate results of the agent's performance on the various datasets can be seen in Table 4, and the trading signals over time in together with the return of the strategies in Figure 3, 4 and 5 for training, validation and out of sample data respectively.

TABLE 4
STRATEGY PERFORMANCE FOR THE DIFFERENT DATA SETS

Data set	Strategy	Relative return [%]	Risk-adjusted return
Training	Buy/hold	20.39	2.40
	Agent	30.86	3.06
Validation	Buy/hold	7.82	1.57
	Agent	13.06	2.13
Out of sample	Buy/hold	-7.20	-2.93
	Agent	2.64	1.89

The returns for each data set are higher for the trading agent than for the simple buy/hold strategy. Adjusting for difference in risk and measuring the risk-adjusted return, the trading agent still outperforms the buy/hold strategy in all examples.

Looking at the results on the out of sample data it seems like the agent is unable to find and capitalize on every trend. The timing of the purchase is when the share price is on its way down and the sale is made just before its peak. Although the agent did not act at the share price peak, or low point trough, it still made money, and more importantly it did not try to trade in the downward trend that appears in the second half of the data. The 2.64% return, realized by the agent is far better than the -7.2% using the buy/hold strategy.

Figure 3
Trading signals and return on training data




Figure 4
Trading signals and return on validation data


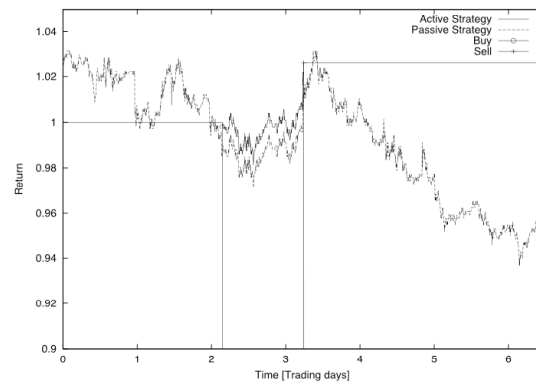


Figure 5
Trading signals and return on out of sample data



The out of sample results highlight an interesting fact, which differs from the previous results. The active agent actually manages to reduce risk. Looking again at Figure 5, the active agent stays liquid during the trend-less first third of the time period. It then trades during a short period and reverts to cash in the negative trend at the last third of the data set. Comparing the two graphs in the figure it actually shows how the simple buy/hold strategy deviates from its mean to a larger extent than the agent's strategy.

V. DISCUSSION

Considering the short time frame of only 6 days for the out of sample validation, it is difficult to draw any solid conclusions about the actual performance of the trading agent. The fact that the agent creates excess risk adjusted returns for the sample period says nothing about its long-run ability to outperform the passive strategy. Even if the agent should manage to out-perform the passive strategy during a longer period of time, the excess return should reflect an increase in risk that comes with active participation in the market, and therefore be consistent with market efficiency.

Despite the set of trading rules, the evolved agent does not, at any time, conduct one bad transaction. Looking at Figure 3, 4 and 5, the agent manages to identify and capitalize on what appears to be trends in the graphs.

The agent evolved from our GP engine trades under a set of stringent rules. It can only buy shares at the ask price and sell shares at the bid price, it is not allowed to either short-sell stocks or go on a margin. *The transaction costs are replicated from an existing broker*, but not necessarily one of the cheapest and the agent is only allowed to buy blocks of shares. The excess cash is left in the broker account without accruing interest. The adopted rules are intended to reflect a conservative trading environment and have the effect of making successful trading increasingly hard.

For both the training, validation and out of sample trading, the agent produces risk adjusted returns higher than the passive buy/hold strategy. The fact that the standard deviation is lower for the trading agent than the passive strategy in the out of sample data is very interesting, considering that the sample was for data with a clear negative trend. If this holds as a general rule for the trading strategy, then it might be so that active trading actually reduces risk in a negative market, while producing higher returns in a market with a positive trend.

VI. FURTHER WORK

The results invite to more research on the use of Genetic Programming in financial markets. Previous research has typically used few variables and a longer perspective in the terminal set, whereas we utilize a lot of variables during a shorter time frame. Combining our way of analyzing multiple variables with a longer time period might produce results that are statistically significant and show positive risk-adjusted returns, greater than those achieved through passive trading strategies.

Assigning more feedback variables increases the ability of the GP system to more easily detect moving averages and other kinds of time lagged information. Multivariate GP systems with a lot of feedback variables should also be well equipped to successfully trade a portfolio of more than one stock, and depending on the definition of fitness, such a system might also be able to cater for different risk preferences.

VII. CONCLUSION

The agent seems to be able to distinguish between small random changes in the Nokia spread and more stable trends. Looking at both the training and validation data the agent finds trends quite easily and also transact at the respective turning points. The agent's ability to find trends is further proved in the out of sample data as the agent stays inactive, holding cash during the last three days, when the market is in a downturn. Alas, the agent does not trade on circumstantial indications but wait for firm trends to arise and the fact that the agent stays liquid during the downturn results in a lower portfolio variance for the trading agent than for the passive strategy, which opens up a new dimension of using the system tailored for a certain risk preference.

VIII. ACKNOWLEDGEMENTS

The Authors would like to thank Anders Liljeqvist, Reuters Sverige AB, SIX AB and the OM Group and NUTEK for their assistance in creating this paper.

IX. REFERENCES

- [1] Koza, J.R. 'Genetic Programming: On the programming of computers by the means of natural selection'. Cambridge, Massachusetts: MIT Press, 1992
- [2] Fama, E.F., "Random Walks in Stock Market Prices," Financial Analysts Journal, September/October 1965
- [3] Fama, E.F., "Market Efficiency, Long-Term Returns, and Behavioral Finance", University of Chicago Date Posted: April 30, 1997, Working Paper Series
- [4] Malkiel B., "Efficient Market Hypothesis", New Palgrave Dictionary of Money and Finance, Macmillan, London, 1992, pp739
- [5] Allen, F., Karjalainen, R., "Using genetic algorithms to find technical trading rules", Working paper at Rodney L. White centre for Financial Research, 1995
- [6] Taylor, M., Allen, H., "The use of technical analysis in the foreign exchange market", Journal of International Money and Finance, 11, 1992, 304-314
- [7] Reilly, F.K., 'Investment Analysis and Portfolio Management', 4th ed, published by Dryden Press, 1994
- [8] Lo, A.W., MacKinlay, A.C., "When are contrarian profits due to stock-market over-reaction?", Review of Financial Studies, 3, 1990, 175-206
- [9] Brock, W., Lakonishok, J., LeBaron, B., "Simple technical trading rules and the stochastic properties of stock returns", Journal of Finance, 47, 1992, 1731-1764
- [10] Campbell, J.Y., Lo, A.W., MacKinlay, A.C., "The econometrics of financial markets", Princeton University Press, 1997
- [11] Tsang E.P.K., Li, J., "Improving technical analysis predictions: An application of Genetic Programming", in proceedings of 12th annual Florida Artificial Intelligence International Research Conference (FLAIRS 99), May 1999
- [12] Tsang, E.P.K., Li, J., "Investment decision making using FGP - A case study", 2000
- [13] Fyfe, C., Marney, J.P., Tarbert, H., "Technical trading versus market efficiency - A Genetic Programming", Applied Financial Economics, 1999
- [14] Marney, J.P., Fyfe, C., Tarbert, H., Miller, D., "Technical analysis versus market efficiency: the quest continues", Barcelona: Computing in economics and finance, 2000
- [15] Iba, H., Sasaki, T., "Financial data prediction by means of Genetic Programming", 2000