

# Improving Technical Analysis Predictions: An Application of Genetic Programming

Jin Li , Edward P.K. Tsang

Department of Computer Science,  
University of Essex, Wivenhoe Park, Colchester,  
United Kingdom  
{jli, edward}@essex.ac.uk

## Abstract

Recent studies in finance domain suggest that technical analysis may have merit to predictability of stock. Technical rules are widely used for market assessment and timing. For example, moving average rules are used to make "buy" or "sell" decisions at each day. In this paper, to explore the potential prediction power of technical analysis, we present a genetic programming based system FGP (Financial Genetic Programming), which specialises in taking some well known technical rules and adapting them to prediction problems. FGP uses the power of genetic programming to generate decision trees through efficient combination of technical rules with self-adjusted thresholds. The generated rules are more suitable for the prediction problem at hand. FGP was tested extensively on historical DJIA (Dow Jones Industrial Average) index data through a specific prediction problem. Preliminary results show that it outperforms commonly used, non-adaptive, individual technical rules with respect to prediction accuracy and average annualised rate of return over two different out-of-sample test periods (three and a half year in each period).

## Introduction

As an approach to financial forecasting, technical analysis is based on the belief that historical price series, trading volume, and other market statistics exhibit regularities. There are two general approaches in technical analysis: one involves qualitative techniques and the other quantitative techniques. The qualitative techniques rely on the interpretation of the form of geometric patterns in the series, such as *double bottoms*, *head-and-shoulders*, and *support and resistance levels*; whilst the quantitative techniques try to create indicators such as *moving average* (MV), *relative strength indicators* (RSI), etc. Notably, both techniques can

be characterised by appropriate sequences of local minima and/or maxima (Neftci 1991).

According to the weak form of the *efficient market hypothesis* (EMH) (Malkiel 1992), since historical price information is already reflected in the present price, technical analysis is useless for predicting future price movements. In recent years, however, this hypothesis has been directly challenged by a fair amount of studies, which supply evidence of predictability of security return from historical price patterns (e.g. Lo & MacKinlay 1990, Brock *et al.* 1992, Campbell *et al.* 1997). The aim of this study is to show how genetic programming (GP) (Koza 1992), a class of algorithms in evolutionary computation, can be employed to improve technical rules. We demonstrate our approach in a particular forecasting task based on the Dow Jones Industrial Average (DJIA).

Quantitative technical rules are often used to generate buy or sell signals based on each rule interpretation. One may want to use technical rules to answer questions such as "*is today a good time to buy if I want to achieve a return of 4% or more within the next 63 trading days?*" and "*is today the right time to sell if I want to avoid a loss of 5% or more within the next 10 days?*" However, the way technical rules are commonly used may not be adequate to answer these questions. How to efficiently apply them and adapt them to these specific prediction problems is a non-trivial task. We propose a GP approach that is capable of combining individual technical rules and adapting the thresholds based on past data. Rules generated by our GP can achieve performances that cannot be achieved by those individual technical rules in their normal usage.

EDDIE (which stands for Evolutionary Dynamic Data Investment Evaluator) is a forecasting system to help investors to make use of the information available to them (Butler 1997, Tsang *et al.* 1998). Such information may include technical rule indicators, individual company's performance indicators, expert predictions, etc. FGP (Financial Genetic Programming) is a descendent of EDDIE. In this paper, we will exam how FGP can be applied to predict whether a return of 4% or more is achievable within the next 63 trading days in the DJIA.

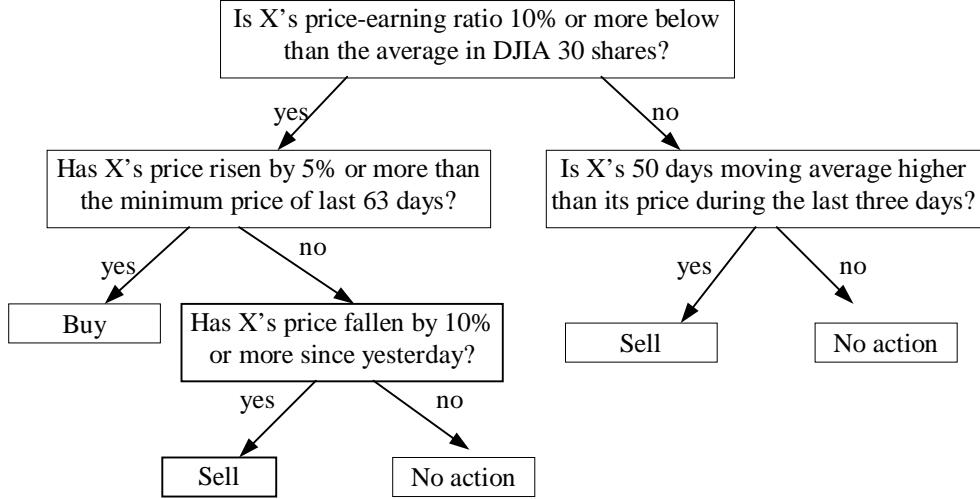


Figure 1. A (simplistic) GDT concerning the actions to take with Share X

## Background of FGP

Genetic programming is a promising variant of genetic algorithms (Holland 1975; Goldberg 1989) that uses tree representations instead of strings. In evolutionary computation, a population (set) of candidate solutions is maintained. For example, a candidate solution could be a decision tree for forecasting. A fitness function is needed to evaluate the quality of each candidate solution with regard to the task to be performed (e.g. how good is a rule for forecasting in our application?). Candidate solutions are selected randomly, biased by their fitness, for involvement in generating members of the next generation. General mechanisms (referred to as *genetic operators*, e.g. reproduction, crossover, mutation) are used to combine or change the selected candidate solutions to generate offspring, which will form the population in the next generation.

Evolutionary computation has been applied to a broad range of problems with some success from traditional optimisation in engineering and operational research to non-traditional areas such as data mining, composition of music (Angeline & Kinnear 1996; Koza 1996) and financial prediction (e.g., Bauer 1994, Mahfoud & Mani 1996, Chen & Yeh 1996, Oussaïdene *et al.* 1997).

In FGP, a candidate solution is represented by a *genetic decision tree* (GDT). The basic elements of GDTs are *rules* and *forecast values*, which correspond to the *functions* and *terminals* in GP. Figure 1 shows an example of a simple GDT. In GP terms, the questions in the example GDT are

*functions*, and the proposed actions are *terminals*, which may also be forecast values.

A GDT can be seen as a set of rules. For example, one of the rules expressed in the GDT in Figure 1 is:

IF X's price earning ratio is 10% or more below than the average in DJIA 30 shares AND X's price has risen by 5% or more than the minimum price of last 63 days, THEN Buy X.

For FGP to work, one must be able to evaluate each GDT. In this paper, we use *prediction accuracy* (the percentage of correct predictions) as fitness function. Our FGP maintains a set of GDTs called a *population* and works in iterations. In each iteration, FGP creates a new generation of population using standard genetic crossover, mutation and reproduction operators. FGP uses *tournament* as its selection strategies.

There are many variations in the way the initial population is generated, the way that the population is updated, the way that crossover and mutation is done, etc. (e.g. see Angeline & Kinnear 1996, Koza *et. al.* 1996). These will not be elaborated here.

## FGP for prediction in DJIA index

We took the Dow Jones Industrial Average (DJIA) index data from 7 April 1969 to 11 October 1976 (1,900 trading days) as training data (or in-sample data) to generate GDTs, and tested them on the data (or out-of-sample data) from 12 October 1976 to 5 May 1980 (900 trading days),

which we shall refer to as "test data I". We used a population size of 1,800, crossover rate of 90%, reproduction rate of 10% and a mutation rate of 1%. The termination condition was 2 hours on a Pentium PC (200 MHz) or 30 generations, whichever reached first. 20 runs were completed in our experiments. Each run generated one GDT. The generated rules were used to predict whether the following goal is achievable at any given day:

*Goal G: the index will rise by 4 % or more within the next 63 trading days (3 months).*

Accordingly, each trading day is classified into "buy" category if G holds or "not-buy" category if G does not hold. The numbers of trading days that belong in each category are roughly same over both the whole training and test period.

We used {If-then-else, And, Or, Not,  $<$ ,  $>$ } as functions. Terminals were conclusions, numbers or indicators. Conclusions could be either *Positive* (meaning that G is predicted to be achievable) or *Negative*. Six technical indicators were derived from rules in the finance literature, such as (Brock *et. al.* 1992; Fama & Blume 1966; Sweeney 1988). They are listed as follows:

- (1)  $MV_{12} = \text{Today's price} - \text{the average price of the previous 12 trading days}$
- (2)  $MV_{50} = \text{Today's price} - \text{the average price of the previous 50 trading days}$
- (3)  $Filter_{-5} = \text{Today's price} - \text{the minimum price of the previous 5 trading days}$
- (4)  $Filter_{-63} = \text{Today's price} - \text{the minimum price of the previous 63 trading days}$
- (5)  $TRB_{-5} = \text{Today's price} - \text{the maximum price of the previous 5 trading days (based on the Trading Range Breakout rule [Brock *et. al.* 1992])}$
- (6)  $TRB_{-50} = \text{Today's price} - \text{the maximum price of the previous 50 trading days}$

Each of the above six indicators is related to some technical analysis rules in the literature. We compared the corresponding six individual technical rules with the GDTs generated by FGP in terms of two criteria: *prediction accuracy* and *average annualised rate of return* (AARR). A recommendation is correct if the goal can be achieved when the recommendation is "buy", or the goal cannot be achieved when the recommendation suggests "do not buy". The *prediction accuracy* of a program or a rule measures the proportion of correct recommendations made by that program or rule. Prediction accuracy is used for evaluating the program's performance because this is what we train the program with. Given any prediction, no matter how accurate they are, the actual return to an investor depends on the investment behaviour, which vary from investor to investor. For reference, we use a simple hypothetical trading behaviour later.

First, we shall explain what our program is compared against. In (Tsang *et al* 1998), FGP was compared with random decisions with a uniformly distributed 50% chance. In this paper, we show that FGP can do better than simple rules that use the input indicators, as it considers the interaction between the indicators. We compare FGP with six individual technical rules that use the above six indicators to generate "buy" or "not-buy" signals in the following ways. The *moving average rules* (1) and (2) generate "buy" signals if today's price is greater than the average price of the preceding  $n$  days (where  $n = 12$  and 50 respectively). The *filter rules* (3) and (4) generate "buy" signals if today's price has risen by 1% or more over the minimum price of the previous  $n$  days ( $n = 5$  and 63 respectively). Here 1% is a threshold that an investor has to choose. The *trading range breakout rules* (5) and (6) generate "buy" signals if today's price is greater than the maximum price over the previous  $n$  days ( $n = 5$  and 50 respectively). AARR was calculated based on the following trading behaviour:

*Hypothetical trading behaviour: we assume that whenever a buy signal is indicated by a rule, one unit of money was invested in a portfolio reflecting the DJIA index. If the DJIA index does rise by 4% or more at day  $t$  within the next 63 days, then we sell the portfolio at the index price of day  $t$ . If not, we sell the portfolio on the 63rd day, regardless of the price. We annualise the return of each unit invested; for example, if 4% is achieved at the 21st trading day (i.e. one month), then the annualised return is  $(4 \times 12 =) 48\%$ . We refer to the mean of these annualised returns as AARR.*

For simplicity, we ignored transaction costs and bid-ask spread. Rules generated by FGP were tested against the above six individual technical rules in the test data. Results are shown in the column of "On test data I" in table 1. Among the six technical rules, Filter\_5 performed best in this set of data. It achieved an accuracy of 52.67% and an AARR of 23.03%. The 20 GDTs achieved an accuracy of 57.97% in average and an average AARR of 27.79%, which is better than the Filter\_5 rule. In fact, even the poorest GDT achieved an accuracy of 53.00% (GDT 18) and AARR of 23.57% (GDT 2), which are still better than the Filter\_5 rule. Our results show conclusively that FGP is capable of generating good rules based on the same indicators used by the technical rules.

To test the robustness of the 20 GDTs across different periods, we applied them to a more recent period, from 10 April 1981 to 29 October 1984 (900 trading days), which we shall refer to as "test data II". The test results are illustrated in the column of "On test data II" in table 1.

20 GDTs	On test data I		On test data II	
	Accuracy	AARR	Accuracy	AARR
GDT 1	60.22%	27.56%	53.89%	55.25%
GDT 2	53.67%	23.57%	53.11%	54.74%
GDT 3	62.00%	31.71%	57.89%	57.80%
GDT 4	58.00%	36.55%	52.33%	61.81%
GDT 5	60.22%	28.23%	63.33%	58.53%
GDT 6	55.11%	29.76%	55.00%	67.00%
GDT 7	61.33%	30.52%	58.00%	56.23%
GDT 8	57.89%	27.16%	54.00%	54.54%
GDT 9	60.67%	28.75%	61.56%	57.69%
GDT 10	55.78%	26.34%	59.11%	59.83%
GDT 11	62.44%	25.93%	55.22%	57.37%
GDT 12	56.78%	25.88%	56.44%	53.19%
GDT 13	56.11%	26.85%	56.44%	59.91%
GDT 14	60.56%	29.66%	54.89%	53.12%
GDT 15	54.78%	25.43%	55.11%	58.05%
GDT 16	56.00%	25.82%	58.11%	59.10%
GDT 17	60.56%	29.18%	58.00%	58.82%
GDT 18	53.00%	23.82%	58.67%	56.61%
GDT 19	60.67%	28.80%	62.89%	58.56%
GDT 20	53.67%	24.18%	57.22%	56.38%
<b>Highest</b>	62.44%	31.71%	63.33%	67.00%
<b>Lowest</b>	53.00%	23.57%	52.33%	53.12%
<b>Mean</b>	<b>57.97%</b>	<b>27.79%</b>	<b>57.06%</b>	<b>57.73%</b>
<b>Standard Deviation</b>	3.07%	3.06%	3.06%	3.15%
<b>6 Technical Rules</b>	(the best result in each column is highlighted)			
MV_12	51.44%	20.68%	44.89%	36.66%
MV_50	42.56%	16.94%	41.89%	46.85%
TRB_5	49.44%	18.18%	47.44%	55.33%
TRB_50	47.44%	-5.34%	48.67%	<b>67.00%</b>
Filter_5 (1%)	<b>52.67%</b>	<b>23.03%</b>	<b>49.44%</b>	54.53%
Filter_63 (1%)	50.56%	22.77%	48.89%	48.76%

Table 1. Performance comparisons between 20 FGP-generated GDTs and six technical rules on test data I (12/10/1976-05/05/1980-900 trading days) and on test data II (10/04/1981-29/10/1984-900 trading days) of the Dow Jones Industrial Average (DJIA) Index

The GDTs achieved an average accuracy of 57.06%, which out-performs all the six technical rules. As in test data set I, even the poorest GDT performed better than all the technical rules on prediction accuracy. The GDTs achieved an average AARR of 57.73%, which is also better than AARRs produced by the technical rules except the

TRB\_50 rule<sup>1</sup>. Test results on data set II further demonstrate the quality of the GDTs generated by FGP.

Two issues are worth pointing out. First, although the number of runs is relatively small, the results are significant because the amount of data tested is large and the re-

<sup>1</sup> Note that the TRB\_50 rule is not particularly reliable. It achieved the lowest AARR in test data I (-5.34%) but the highest AARR in test data II (67.00%). The erratic performance of the TRB\_50 rule is partly due to the fact that it generates very few buy signals.

sults are consistent. It is encouraging to see that our GDTs achieve nearly the same mean of accuracy (57.97%, 57.06%) with almost the same standard deviation (3.07%, 3.06%) over two test periods. Second, our calculation of AARR assumes that funds are always available whenever a positive position is predicted, and such funds have no cost when idle. Exactly how one can make use of the predictions by the GDTs is beyond this paper.

## Conclusion and further work

It is not our role to defend technical analysis here, although our results show that there is some predictability in the DJIA index based on historical data alone. Our main objective is to illustrate that by taking indicators used in technical rules as input, FGP, a genetic programming based system, can generate decision trees that perform better than those technical rules. For the specific task tested, FGP reliably generated accurate GDTs that perform better than the individual technical rules. This involves combining indicators in individual technical rules and finding thresholds in different parts of the decision trees.

The application presented here is not complete since important issues such as transaction costs and capital adequacy were ignored. In the future, we plan to consider these factors. We intend to bring in constraint satisfaction techniques, which have been demonstrated to be useful in genetic algorithms (Lau & Tsang 1997; Tsang 1993).

## Acknowledgements

This work is partly supported by the Research Promotion Fund, 1997 (DDP540), University of Essex. Jin Li is supported by an Overseas Research Students Award (ORS) and a University of Essex Studentship. The DJIA index data was generously provided by Professor Blake LeBaron at University of Wisconsin - Madison. The anonymous referees provided valuable comments.

## References

Angeline, P. & Kinnear, K.E., eds. (1996). *Advances in genetic programming II*, MIT Press.

Bauer, R. J. Jr. (1994). *Genetic Algorithms and Investment Strategies*. New York, John Wiley & Sons, Inc.

Brock, W.; Lakonishok, J. and LeBaron, B. (1992). Simple technical trading rules and the stochastic properties of stock returns. *Journal of Finance*, 47, 1731-1764.

Butler, J. M. (1997). Eddie beats the market, data mining and decision support through genetic programming. *Developments*, Reuters Limited, Vol1.

Campbell, J. Y.; Lo, A. W.; & MacKinlay, A. C. (1997). *The econometrics of financial markets*. Princeton, N.J.: Princeton University Press.

Chen, S-H & Yeh, C-H. (1996). Toward a computable approach to the efficient market hypothesis: An application of genetic programming, *Journal of Economic Dynamics and Control*, 21, 1043-1063.

Fama, E. F. & Blume, M. E. (1966). filter rules and stock-market trading. *Journal of Business* 39(1), 226-241.

Goldberg, D.E. (1989). *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley.

Holland, J. H. (1975). *Adaptation in natural and artificial system*. University of Michigan Press.

Koza, J. R. (1992). *Genetic Programming: on the programming of computers by means of natural selection*. MIT Press.

Koza, J.; Goldberg, D.; Fogel, D. and Riolo, R. eds. (1996). *Proceedings of the First Annual Conference on Genetic programming*. MIT Press.

Lau, T. L. & Tsang, E.P.K. (1997). Solving the processor configuration problem with a mutation-based genetic algorithm. *International Journal on Artificial Intelligence Tools* (IJAIT), World Scientific, Vol.6, No.4, 567-585.

Lo, A.W. & MacKinlay, A. C. (1990). When are contrarian profits due to stock market overreaction? *Review of Financial Studies* 3, 175-206.

Mahfoud, S. & Mani, G. (1996). Financial Forecasting Using Genetic Algorithms. *Journal of Applied Artificial Intelligence* Vol.10, Num 6, 543-565.

Malkiel, B. (1992). Efficient market Hypothesis. In New- man, P.; Milgate, M. and Eatwell, J.(eds.), *New Palgrave Dictionary of Money and Finance*, Macmillan, London, pp739.

Neftci, S. N. (1991). Naïve trading rules in financial markets and Wiener-Kolmogorov prediction theory: A study of 'technical analysis'. *Journal of Business*, 64, 549-571.

Oussaidene, M.; Chopard, B.; Pictet, O. and Tomassini, M. (1997). Practical aspects and experiences - Parallel genetic programming and its application to trading model induction. *Journal of Parallel Computing* Vol. 23, No. 8, 1183-1198.

Sweeney, R. J. (1988). Some new filter rule tests: Methods and results. *Journal of Financial and Quantitative Analysis*, 23, 285-300.

Tsang, E. P. K.; Butler, J. M. & Li, J. (1998). EDDIE beats the bookies. *International Journal of Software, Practice & Experience*, Wiley, Vol.28 (10), 1033-1043.

Tsang, E.P.K. (1993). *Foundations of constraint satisfaction*. Academic Press, London.