ELSEVIER

EDDIE-Automation, a decision support tool for financial forecasting

Edward Tsang*, Paul Yung, Jin Li

Department of Computer Science, University of Essex, Colchester CO4 3SQ, UK

Available online 2 July 2003

Abstract

Evolutionary Dynamic Data Investment Evaluator (EDDIE) is a genetic programming (GP)-based decision support tool for financial forecasting. EDDIE itself does not replace forecasting experts. It serves to improve the productivity of experts in searching the space of decision trees, with the aim to improve the odds in its user's favour. The efficacy of EDDIE has been reported in the literature. However, discovering patterns in historical data is only the first step towards building a practical financial forecasting tool. Data preparation, rules organization and application are all important issues. This paper describes an architecture that embeds EDDIE for learning from and monitoring the stock market.
© 2003 Elsevier B.V. All rights reserved.

Keywords: Financial forecasting tools; Genetic programming

1. Background

This paper describes a decision support tool for financial forecasting. Suppose, through diligent research, one discovers that a particular ratio, alongside with other commonly known factors (such as interest rate, money supply, dividends, moving average and volatility), is highly relevant to the movement of a share, but do not know the exact mathematical relation. The question is: could this discovery give this person any advantage in trading? In other words, how could one take advantage of such knowledge? The answer is: discovering relevant factors alone does not offer too much help. What needs to be done

further is nontrivial. Even if one limits oneself to using decision trees as in ID3/C4.5/See5/C5.0 [17,18] or genetic programming [7,10,11], one still needs to search in a huge space of possible decision trees. The search space is exponential in size to the maximum depth that one may want to look at and it is quite impossible to search exhaustively for useful trees, even with the help of fast computers at today's standard [11].

EDDIE (which stands for *Evolutionary Dynamic Data Investment Evaluator*) is an interactive tool, based on genetic programming (GP), to help analysts to search the space of interactions between factors (in the form of decision trees) and make financial decisions. As the aim of this paper is to describe the role of EDDIE in a decision support system (as opposed to describing its forecasting ability, which has been reported in Refs. [13, 14,24], and will be briefly summarized below), a

---

* Corresponding author.
E-mail address: edward@essex.ac.uk (E. Tsang).
URL: http://cswww.essex.ac.uk/CSP/finance.

survey of scientific financial forecasting methods will not be presented (see Refs. [3,7,8,14,19] for a few examples) here.

The contribution of EDDIE is effectively searching for combinations (interactions) of financial indicators and discovering thresholds (to be elaborated in the next section). However, after developing a system that can mine patterns in the data, many issues are left to be studied in building a practical financial forecasting tool. For example, one must study *how* such predictions, which are at best statistically reliable (i.e. there is no guarantee that they will be correct every time), for investment. This is *not* the scope of this paper. In this paper, we investigate the practical issues in generating and using such predictions. In the remaining part of the paper, we shall describe an architecture that employs EDDIE in learning from and monitoring the stock market. We shall argue that the implemented automation system enables the user to conduct large-scale learning and monitoring in the stock market efficiently.

2. EDDIE for financial forecasting

EDDIE is a decision support tool based on genetic programming [2,9–11]. The efficacy of EDDIE has been reported in Refs. [5,12,14,21–24]. A brief summary is given in this section.

The first implementation of EDDIE, called EDDIE-1, was applied to horse racing [22]. Its performance was promising but limited, due to the lack of data in horse racing at the time. FGP-1 (co-named EDDIE-3; FGP stands for *Financial GP*), an implementation of EDDIE, was applied to financial forecasting [13]. One of the main applications of EDDIE was to generate and test hypotheses of the following form:

Will the price of share/index X rise by *r*% within the next n trading days?[1]

---

[1] In fact, one can ask whether prices will *fall* by *r*% within the next *n* days, or, more generally, within the next *n* units of time where each row in the data set represents information at each unit of time. In Ref. [15], EDDIE was used to monitor arbitrage opportunities by finding patterns in tick data.

For example, one may want to investigate whether the price of the FTSE-100 index will rise by 2.2% within the next month. The syntax of the trees to be generated is as in Fig. 1.

A "Positive" prediction means that the goal can be achieved; "Negative" means otherwise. An example decision tree (annotated for readability) is shown below:

$$(\text{IF}(\text{MV}\_50 < -18.45) \text{ THEN Positive}$$
$$\text{ELSE } (\text{IF}((\text{TRB}\_5 > -19.48)$$
$$\text{AND } (\text{Filter}\_63 < 36.24)) \text{ THEN Negative}$$
$$\text{ELSE Positive}))$$

This decision tree actually represents three simple rules: (a) if MV_50 is less than -18.45, then positive is predicted; (b) if MV_50 is not less than -18.45, and TRB_5 is greater than -19.48, and Filter_63 is less than 36.24, then negative is predicted; and (c) otherwise (not elaborated here) positive is predicted.

In GP terms, FGP-1 uses {If-then-else, And, Or, Not, $<$, $=$, $>$} as functions and indicators, numbers and predictions as terminals. The crossover operator was designed to take care of the type of the branches. Indicators were taken from the finance literature, such as Refs. [1,4,6,18,20], normalized. For example, MV_12 is the 12-day moving average divided by the mean in the last 12 days. Vol_50 is the standard deviation of the prices in the last 50 days divided by the last 50 days' mean. Details of the parameters, which are not central to the ideas in this paper, can be found in Ref. [12].

EDDIE is useful to its users because even if the indicators are relevant to the prediction of the prices, finding (a) the logical interaction between the various indicators and (b) the thresholds (e.g. -18.45 in the above example) could be extremely laborious without tools such as genetic programming and neural networks.[2] There is no magic behind genetic programming or neural networks. All they do is to efficiently explore the space of possible patterns.

FGP-2 (co-named EDDIE-4) extended FGP-1 by using a constrained fitness function, which allows the user to adjust the level of cautiousness [24]. When the user instructs FGP-1 to recommend cautious rules, FGP-2 will attempt to improve precision, possibly at

---

[2] Neural networks are perfectly capable of finding nonlinear patterns [7], though they are considered by some as black-boxes.

```
        Tree  :=  "If-then-else" <Condition> <Tree> <Tree> | <Prediction>
  <Condition>  :=  <Condition> "And" <Condition>  |  <Condition> "Or" <Condition>  |
                    "Not" <Condition>  |  <Indicator> <RelationOperation> <Threshold>
  <Indicator>  :=  "MV_12" | "MV_50 " | "Filter_5" | "Filter_63"| "TRB_5 " | "TRB_50" |
                    "Vol_12" | "Vol_50
<RelationOperation>  :=  ">"  |  "<"  |  "="
  <Threshold>  :=  Number
 <Prediction>  :=  "Positive" | "Negative"
```

Fig. 1. The BNF grammar used by FGP.

the price of missing opportunities. This means FGP-2 will make fewer recommendations to buy (or sell). The effectiveness of FGP-2 as a forecasting tool has been reported in Refs. [12,14,24]. FGP-2 is designed as a forecasting tool, in which the user is given control over the percentage of opportunities to pick up [24] (as opposed to incorporating a *misclassification cost* [18] in the objective function).

3. EDDIE-Learn 1.0, an automated learning architecture

Finding patterns in a given set of data is not the only laborious task in forecasting. The automation of the learning process is a necessary step towards turning EDDIE into an effective forecasting support tool. Here are some of the reasons.

1. Data preparation plays an important part in data mining. One does not always know which indicators are relevant to the movement of a particular share. One often has to prepare different sets of indicators and data-mine in them.
2. One often needs to experiment with a wide variety of hypotheses in order to feel one's way in the data.
3. To apply GP effectively, one has to experiment with different parameter sets, such as the population size, number of iterations and level of cautiousness.

For example, over a dozen versions of data sets and various hypotheses have been experimented and a number of hypotheses have been tested when we applied EDDIE to finding arbitrage opportunities [14]. Preparation of data took hours of laborious manual and programming work.

To serve the above needs, an EDDIE-based automatic learning system has been implemented. Input to the system includes:

(a) Data files, in comma-separated values (CSV) format; in the implemented system, this includes the date and share prices. It is possible to include new attributes in the data;
(b) The range file, in text format, which defines the range of data to be used for training and testing;
(c) The control file, in text format, which defines the GP parameters, including population size, mutation rate, etc.;
(d) Program parameters: the rate of return (r), forecasting horizon (n) and the precision threshold (p). Together, they define the target return (r% within the next n days) and the minimum precision (on the test data) for the rules to be accepted.

For each run, a report in CSV format (which is suitable for spreadsheet use as well as electronic processing) is generated. The main output of EDDIE-Automation 1.0 is a collection of rules with precision above the threshold specified (default is 75%). Intermediate data files are generated for verification purpose.

Fig. 2 gives an overview of the EDDIE-based automatic learning system, EDDIE-Learn 1.0. To use the system, the user stores the data files of all the shares that he/she wants EDDIE to learn from in the directory dir_input. Each data file contains a time series, in the form of time (e.g. dates) and price pairs. The program autofeedfgp completes the rest of the training by invoking a sequence of programs: The program autofeed calculates the indicators for each file in dir_input—after adding extra columns to the
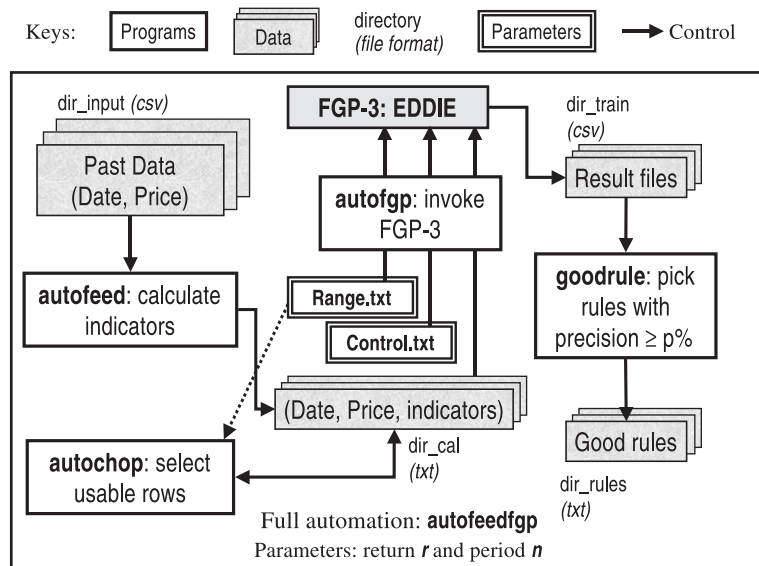
Fig. 2. EDDIE-Learn 1.0, the Learning Architecture that embeds EDDIE.

table, the augmented data will be stored in the directory dir_cal. The program autochop selects the rows specified in the parameter file range.txt (explained in point (b) above). The data prepared is fed into an automation of EDDIE, FGP-3. FGP-3 also reads in two parameter files, range.txt and control.txt (explained in point (c) above).

Training results are stored in a directory (called dir_train) in CSV format (to enable users to inspect them with a spreadsheet). The program goodrule picks the rules which performance is above the specified precision threshold p (explained in point (d) above), and put them in a good-rules bank (dir_rules). Care has been taken in generating the names of the files, which reflect the return (r), period (n) and precision (p). Multiple rules may be generated for the same share—this leaves us with the option to combine rules and recommendations in the future.

4. EDDIE-Automation 1.0, an automated monitoring architecture

In the previous section, we have explained how EDDIE-Learn 1.0 learns patterns and put them into a good-rules bank. In this section, we present EDDIE-Automation 1.0, an architecture that automates the

whole learning and monitoring process. The overview of the architecture is shown in Fig. 3.

In order to apply the rules learned, a Rules Application Program (RAP) is implemented. Given (a) a decision tree DT and (b) a record R, which comprises the price and the indicators that DT used in its training, RAP returns "True" if R indicates a positive position according to DT (meaning that r% can be achieved within the next n days for the r and n specified in the parameter file range.txt), and "False" otherwise.

The decision trees in the good-rules bank are used to monitor their corresponding shares in order to detect investment opportunities. To facilitate that, live data is fed into the directory dir_live in Fig. 3. The program autofeed (which is also used in EDDIE-Learn 1.0, see Fig. 2) prepares the data for RAP. This will provide RAP with a record with the prices and other indicators necessary for its calculation. The program autorap picks up each rule in the good-rules bank in turn, and apply it to the corresponding record prepared by autofeed (correspondence is established through their file names). The results of all predicted investment opportunities are stored in a directory called dir_alert. All investment opportunities in dir_alert are collected and summarized in a text file, which can be delivered to investors by email or text messages by phone if desired.
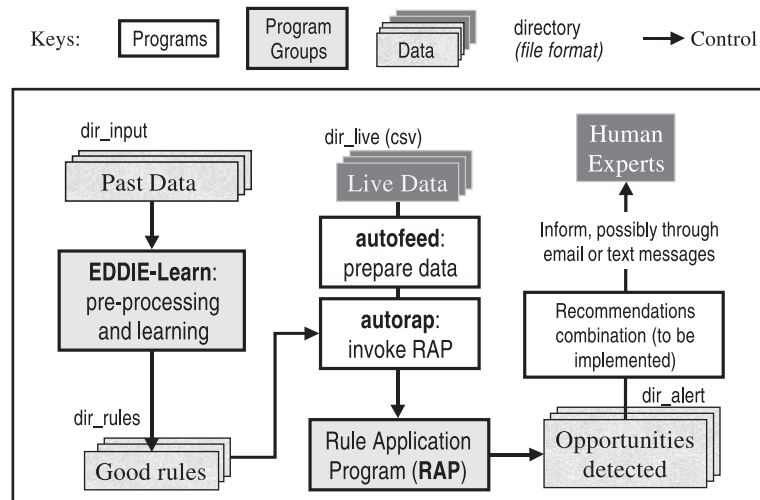
Fig. 3. The EDDIE-Automation Architecture.

In the good-rules bank, multiple rules may apply to the same share. Recommendations may or may not agree with each other. In the future, we would consider combining recommendations before communicating with the users. There are many ways to do so. For example, one may only decide to buy a particular share if at least m out of the n rules in the good-rules bank make positive recommendations. Or, one may like to generate rules for both long-term and short-term trading, and only act if both long-term and short-term prospects are positive. Different ways of combining rules are on the agenda of the EDDIE project.

EDDIE-Automation 1.0 was built in modules. This allows flexibility in the system. For example, should the input files take a different format, or should the user require a different set of indicators, only the autofeed module needs to be changed. Even the core module FGP-3 could be replaced by another learning system that uses the same input and output format (such as C4.5 [17] and See5/C5.0 [18]).

All parts of EDDIE-Automation 1.0 shown in Fig. 3 have been implemented except live data-feed, combining recommendations and communication with the user. In the system implemented, live data is placed into a directory (dir_live in Fig. 3) manually. FGP-3 was implemented in C++. RAP was implemented in Java. The rest of the programs were written in Perl script. All experiments were run under Windows 2000.

5. EDDIE-Automation 1.0 as a decision support system

First, it is worth noting that although the implementation generates technical rules only, the EDDIE-Learn architecture is not limited to technical analysis (although there are good arguments for studying technical analysis scientifically [6,16]). One could also use indicators that are generated by economic models. This was done in EDDIE-ARB, where the theoretical price of options was also used for detecting arbitrage opportunities [14].

Second, it is worth re-iterating that EDDIE-Automation is no replacement for experts. EDDIE does not guarantee 100% correctness in its predictions. Instead, it aims to improve the odds in its user's favour. It helps to find patterns in past data, which can be used for reference by its user. It cannot predict new events or market crashes, unless its user can. (In general, EDDIE has no prediction power unless the user does.) What EDDIE can do is to significantly enhance the user's productivity in finding patterns and monitoring the market. As a decision support tool, EDDIE-Automation should be judged by how and how much it could improve the user's productivity. We shall therefore look at various aspects of financial data mining below.

Data preparation is crucial to data mining. It is often labour-intensive too. One often needs to experiment with different ways to create new attributes (such as

new ratios), remove attributes (in order to reduce the search space), remove records (to reduce noise or reduce complexity), etc. before feeding data into any learning system. Data preparation often involves domain knowledge. While EDDIE-Learn 1.0 cannot offer much help in this aspect, it provides a framework to automate data preparation once the user has decided what indicators to create, which columns or rows to use for training and testing, etc. By supplying EDDIE-Learn 1.0 with an appropriate module to prepare tables from raw data, users can experiment with a large number of data sets automatically.

Where EDDIE-Learn 1.0 can provide significant help to its user is in hypotheses testing. Users may ask questions of the format "will any of these shares rise by r% within the next n days?". After the user stores the price series of all the shares of interest into the input directory (dir_input in Fig. 2), a simple instruction will set EDDIE-Learn 1.0 off to learn rules for all the shares in that directory. Users may easily repeat the experiments by varying r and n.

Going back to the points made in Section 1, EDDIE helps the users to discover interactions between financial indicators in the form of decision tree. Each branch in the decision tree represents the combination of a number of indicators. It may be the case that indicator X is only relevant if indicator Y is also considered. EDDIE will also search for threshold values. Given the exponential size of search space, it is in practice impossible for human users to discover such thresholds manually. There is no guarantee that EDDIE will find patterns, but it stands a better chance to do so than its human user as it can work day and night continuously.

When FGP-3 is asked to act cautiously, each decision tree in the good-rules bank may pick up only a small number of investment opportunities. This is not a serious problem, as one can use EDDIE-Learn 1.0 to learn as many rules as one's computation power permits—with the hope (but no guarantee) that different opportunities will be picked up by different rules. With a large number of rules in the good-rules bank, it is necessary to automate the monitoring process. EDDIE-Automation 1.0 provides a framework to monitor the market and report to the users when opportunities are detected.

Obviously, the more shares and indices that one studies, the more investment opportunities one could potentially discover. Besides, the more rule sets one generates, the more chance that one could spot investment opportunities or combine recommendations. EDDIE-Automation 1.0 enables the users to conduct large-scale learning to fill the good-rules bank (directory dir_rules in Figs. 2 and 3). EDDIE-Automation 1.0 is a useful tool because the amount of work involved, in both learning and monitoring, is huge and therefore beyond human effort.

To summarize, this paper describes the role of EDDIE in a decision support system. EDDIE enables the user to discover patterns that he/she could not practically discover manually. EDDIE-Learn 1.0 enables the user to do large-scale learning. EDDIE-Automation 1.0 is a practical tool for forecasting. Finally, it is worth re-iterating that EDDIE is only a useful decision support tool, not a replacement for forecasting experts.

Acknowledgements

References

[1] S.S. Alexander, Price movement in speculative markets: trend or random walks, No. 2, in: P. Cootner (Ed.), The Random Character of Stock Market Prices, MIT Press, Cambridge, MA, 1964, pp. 338–372.

[2] P. Angeline, K.E. Kinnear Jr. (Eds.), Advances in Genetic Programming II, MIT Press, 1996.

[3] R.J. Bauer Jr., Genetic Algorithms and Investment Strategies, Wiley, 1994.

[4] W. Brock, J. Lakonishok, B. LeBaron, Simple technical trading rules and the stochastic properties of stock returns, Journal of Finance 47 (1992) 1731–1764.

[5] J.M. Butler, EDDIE Beats the Market, Data Mining and De-

cision Support Through Genetic Programming, Developments, vol. 1, Reuters, 1997 July.

[6] E.F. Fama, M.E. Blume, Filter rules and stock-market trading, Journal of Business 39 (1) (1966) 226–241.

[7] S. Goonatilake, P. Treleaven (Eds.), Intelligent Systems for Finance and Business, Wiley, New York, 1995.

[8] N.K. Kasabov, Foundations of Neural Networks, Fuzzy Systems and Knowledge Engineering, MIT Press, 1996.

[9] J.R. Koza, Genetic Programming: On The Programming of Computers by Means of Natural Selection, MIT Press, 1992.

[10] J. Koza, D. Goldberg, D. Fogel, R. Riolo (Eds.), Proceedings, First Annual Conference on Genetic Programming, MIT Press, 1996.

[11] W.B. Langdon, R. Poli, Foundations of Genetic Programming, Springer, 2001.

[12] J. Li, FGP: a genetic programming based tool for financial forecasting, PhD Thesis, University of Essex, UK, 2001.

[13] J. Li, E.P.K. Tsang, Investment decision making using FGP: a case study, Proceedings of Congress on Evolutionary Computation (CEC'99), Washington DC, USA, July 6–9, 1999.

[14] S. Mahfoud, G. Mani, Financial forecasting using genetic algorithms, Applied Artificial Intelligence 10 (1996) 543–565.

[15] S. Markose, E. Tsang, H. Er, EDDIE for stock index options and futures arbitrage, in: S.-H. Chen (Ed.), Genetic Algorithms and Genetic Programming in Computational Finance, Kluwer Academic Publishing, 2002, pp. 281–308.

[16] C. Neely, P. Weller, R. Ditmar, Is technical analysis in the foreign exchange market profitable? A genetic programming approach, in: C. Dunis, B. Rustem (Eds.), Proceedings, Forecasting Financial Markets: Advances for Exchange Rates, Interest Rates and Asset Management, London, May, 1997.

[17] J.R. Quinlan, C4.5: Programs for Machine Learning, Morgan Kaufmann, San Mateo, 1993.

[18] J.R. Quinlan, Data mining tools See5 and C5.0, http://www.rulequest.com/see5-info.html (accessed 3 February 2003).

[19] D. Sornette, W.-X. Zhou, The US 2000–2002 market descent: how much longer and deeper? Quantitative Finance 2 (6) (2002) 468–481.

[20] R.J. Sweeney, Some new filter rule test: methods and results, Journal of Financial and Quantitative Analysis 23 (1988) 285–300.

[21] E.P.K. Tsang, J. Li, Combining ordinal financial predictions with genetic programming, Proceedings, Second International Conference on Intelligent Data Engineering and Automated Learning, Hong Kong, December 13–15, 2000.

[22] E.P.K. Tsang, J. Li, J.M. Butler, EDDIE beats the bookies, International Journal of Software, Practice and Experience, vol. 28 (10), Wiley, 1998, pp. 1033–1043.

[23] E.P.K. Tsang, J. Li, S. Markose, H. Er, A. Salhi, G. Iori, EDDIE in financial decision making, Journal of Management and Economics 4 (4) (2000 November) (http://www.econ.uba.ar/www/servicios/publicaciones/journal4/Index.htm).

[24] E.P.K. Tsang, J. Li, EDDIE for financial forecasting, in: S.-H. Chen (Ed.), Genetic Algorithms and Programming in Computational Finance, Kluwer Academic Publishing, 2002, pp. 161–174.

Dr. Edward Tsang is a professor in computer science at the University of Essex. He specializes in applied artificial intelligence, especially heuristic search, stochastic methods and their application to constraint satisfaction.

Paul Yung received his Master degree from the Department of Electronic Systems Engineering, University of Essex. He is a freelance programmer specialized in telecommunication and Unix systems.

Dr. Jin Li received his PhD in Computer Science from University of Essex. He is a freelance programmer specialized in Window and web-based software design.