Expert Systems With Applications 54 (2016) 193–207

Contents lists available at ScienceDirect

Expert Systems With Applications

journal homepage: www.elsevier.com/locate/eswa

Evaluating machine learning classification for financial trading: An empirical approach

Eduardo A. Gerlein[a,c,*], Martin McGinnity[a,b], Ammar Belatreche[a], Sonya Coleman[a]

[a] Intelligent Systems Research Centre, University of Ulster, Londonderry, UK
[b] School of Science and Technology, Nottingham Trent University, Nottingham, UK
[c] Electronics Department, Pontificia Universidad Javeriana, Bogotá, Colombia

ARTICLE INFO

Keywords:
Trading
Financial forecasting
Computer intelligence
Data mining
Machine learning
FOREX markets

ABSTRACT

Technical and quantitative analysis in financial trading use mathematical and statistical tools to help investors decide on the optimum moment to initiate and close orders. While these traditional approaches have served their purpose to some extent, new techniques arising from the field of computational intelligence such as machine learning and data mining have emerged to analyse financial information. While the main financial engineering research has focused on complex computational models such as Neural Networks and Support Vector Machines, there are also simpler models that have demonstrated their usefulness in applications other than financial trading, and are worth considering to determine their advantages and inherent limitations when used as trading analysis tools. This paper analyses the role of simple machine learning models to achieve profitable trading through a series of trading simulations in the FOREX market. It assesses the performance of the models and how particular setups of the models produce systematic and consistent predictions for profitable trading. Due to the inherent complexities of financial time series the role of attribute selection, periodic retraining and training set size are discussed in order to obtain a combination of those parameters not only capable of generating positive cumulative returns for each one of the machine learning models but also to demonstrate how simple algorithms traditionally precluded from financial forecasting for trading applications presents similar performances as their more complex counterparts. The paper discusses how a combination of attributes in addition to technical indicators that has been used as inputs of the machine learning-based predictors such as price related features, seasonality features and lagged values used in classical time series analysis are used to enhance the classification capabilities that impacts directly into the final profitability.

1. Introduction

Data mining is the process of finding hidden patterns within data using automatic or semi-automatic methods (Witten, Frank, & Hall, 2011). In particular, machine learning (ML) techniques have shown impressive performance in solving real life classification problems in many different areas such as communications (Di, 2007), internet traffic analysis (Nguyen & Armitage, 2008), medical imaging (Wernick, Yang, Brankov, Yourganov, & Strother, 2010), astronomy (Freed & Lee, 2013), document analysis (Khan, Baharudin, Khan, & E-Malik, 2009), biology (Zamani & Kremer, 2011) and time series analysis (Qi & Zhang, 2008). Although complex models such as Neural Networks (NN) and Support Vector Machine (SVM) techniques are studied within the ML field, several other approaches also exist, characterized by a greater degree of simplicity when compared with NN and SVM. Despite this apparent simplicity, some of the ML techniques may be well-suited for quantitative analysis within the financial industry, as their capabilities for finding hidden patterns in large amounts of data may help in financial forecasting for trading.

Financial trading of securities using technical and quantitative analysis has been traditionally modelled by statistical techniques for time series analysis such as the ARMA (Box et al., 1994) and ARIMA models, and the more sophisticated ARCH technique (Engle, 1982). In contrast to these statistical approaches, complex models coming from the ML field have emerged attempting to predict future movements of securities' prices (Yoo, Kim, & Jan, 2007). The extensive literature has shown how some ML techniques specializing in classification and regression tasks have demonstrated

* Corresponding author at: Electronics Department, Pontificia Universidad Javeriana, Bogotá, Colombia. Tel. +5713208320 ext 5549.
E-mail addresses: gerlein-e@email.ulster.ac.uk (E.A. Gerlein), tm.mcginnity@ulster.ac.uk (M. McGinnity), a.belatreche@ulster.ac.uk (A. Belatreche), sa.coleman@ulster.ac.uk (S. Coleman).

http://dx.doi.org/10.1016/j.eswa.2016.01.018
0957-4174/

that they are well-suited for quantitative analysis within the financial industry, as their capabilities of finding hidden patterns in large amounts of financial data may help in derivatives pricing, risk management and financial forecasting. One of the most published projects that uses such techniques in financial applications is Standard & Poor's Neural Fair Value 25 portfolio (Smicklas, 2008), which selects on a weekly basis 25 stocks using an artificial NN, from a total of 3000 stocks yield relative to that of the S&P 500 index, attempting to outperform the market by calculating a stock's weekly fair value based on fundamental analysis. Particularly for securities trading, the utility of complex models such as NN, SVM and hybrid models (Cai, Hu, & Lin, 2012) have been extensively studied and have led to promising results. Nevertheless, information regarding the incorporation of such methods into trading floor operations tends to remain hidden to the public, for commercial proprietary reasons (Yamazaki & Ozasa, n.d.; Duhigg, 2006; Patterson, 2016).

In terms of financial trading, analysts in the industry (usually referred to as "quants") have developed *technical indicators* that are used to identify the most suitable moments to open and close trades, and are possibly the most popular tools currently used in technical trading. Published research that aims to incorporate Computational Intelligence (CI) in financial prediction shows how those *technical indicators* have been used as inputs to ML models to find the hidden patterns or relationships among them, in order to predict future prices, trends or a percentage of confidence in those predictions. With the possible exception of long term averages, those *technical indicators* are constructed using information of prices over short periods in the past, no more than 20–30 trading periods in order to incorporate the historical behaviour in a single value. The selected trading period is part of the trading strategy and might vary from long frames of 1 day to small frames of 1 minute, or even smaller time windows as in the case of high frequency trading. The construction of such indicators can be seen as a process used in large scale time series data sets called *dimension reduction* (Wang et al., 2005) that attempts to transform the series to another domain seeking a version of the data set that might be much simpler to analyse. In contrast to time series analysis where the data set is seen as a whole entity, ML classification tasks construct independent *instances* that are representative examples of the concept to be learned.

Financial predictions that incorporate ML approaches construct the training, test and off-sample data sets as a collection of *instances* using popular *technical indicators* as reported in a number of papers. Hence, an *instance* is created usually using the value of the current price and the instantaneous values of the mentioned indicators, generating a static picture of the situation of the market for the exact time that the *instance* is constructed. In this scenario, each *instance*, i.e. prices and their correspondent *technical indicators* used as attributes, becomes itself an independent example of the problem, which avoids the time dependence in the series, approaching the problem as simple classification task rather than a time series analysis in the strict meaning of the word. The hypothesis in this case is that once a ML model is trained, it may be able to classify individual *instances* using the technical indicators as attributes, due to the fact that those unseen instances represent in turn the invariant circumstances of the market at certain points in time, and that the result of the classification task can be interpreted as a trend forecast. The main implication of this hypothesis is that the financial forecasting can benefit from the use of simpler ML techniques rather than using complex time series analysis approaches, simplifying the use of computational resources while at the same time avoiding indexing and ordering issues in the data sets.

This paper addresses the question of the usefulness of low-complex ML classifiers in financial trading, and in particular will demonstrate if such low-complexity binary classification approaches are able to generate consistent profitable trading over an extensive period of time. The paper's main contribution resides in the fact that simple machine learning models that traditionally have been precluded from financial applications, as opposed to the more complex NN and SVM, can be used to generate profitable transactions on the long term with the correct combination of periodic retraining, training set size and attribute selection. The work is motivated mainly by the results reported in (Barbosa, 2011), which claims outstanding financial results using simple classifiers. In this case, a simple model is characterized by the low computational requirements for both the training and the classifying process due to the inherent simplicity of the learned model (instance-based classifiers, decision trees and rule-based learners). While the main objective of ML classification is to maximise accuracy, this might not be the best metric to evaluate the performance of such systems when used in the context of financial trading. The most important metric when assessing trading strategies is undoubtedly profitability, reflected in this paper as cumulative return over a specific trading period. In this paper, it is developed an empirical comparison between average accuracy and cumulative return as the main metrics of performance of a set of six machine learning models (OneR, C4.5, JRip, Logistic Model Tree, KStar and Naïve Bayes). The models produce a binary classification used later to predict price movement (up or down in the next trading period) for the USDJPY currency pair using six hour time frames over a trading time-frame of six years. The six hour time frame was selected to be able to validate and compare with the results reported in (Barbosa & Belo, 2008b), although the same approach used in the experiments can be applied to different time frames. A set of experiments was conducted where the results of modifying three variables were studied: training set size, period of retraining and number of attributes for the training and test sets. The results show relative low accuracy, only a few points over 50%, but at the same time, very promising results in terms of profitability. Later, further experiments are conducted on simulated trades over the same period of time using EURGPB and EURUSD currency pairs, and similar results are reported.

The remainder of the paper is organised as follows: Section 2 discusses related work using ML in financial forecasting applications. Section 3 presents the general experimental setup, describing the data sets, and the attribute selection to feed the models and briefly describes the different ML algorithms used in the experiments as well as their particular parameter set up in order to present comprehensive information for future experiment replication. Section 4 discusses the results detailing each one of the simulated trading scenarios. Finally, Section 5 concludes the paper and explores future work.

2. Machine learning in financial forecasting

Within the financial trading chain two main areas can be identified, where the use of ML techniques have reported particularly successful implementations: derivatives pricing, risk management and financial forecasting. Financial forecasting is possibly the most important application within ML for data mining in Capital Markets. ML techniques for forecasting include expert or rule-based systems, decision trees, NNs and genetic computing. Applications within the trading cycle such as Algorithmic Trading Engines[1] and

[1] Algorithmic trading engines in the buy side, are essentially semi-automatic computer aided systems that help retail investors to take the best financial decisions in terms of high returns at lowest possible risks, and by means of programming specific rules the system are capable of transmitting pre- and post-trade data about quotes and trades to other market participants (Hendershott, 2003 Chan, 2008). The literature also reports the use of algorithmic trading engines in the sell