Artificial Neural Networks in Financial Trading

A

Bruce Vanstone
Bond University, Australia

Clarence Tan
Bond University, Australia

INTRODUCTION

Soft computing represents that area of computing adapted from the physical sciences. Artificial intelligence (AI) techniques within this realm attempt to solve problems by applying physical laws and processes. This style of computing is particularly tolerant of imprecision and uncertainty, making the approach attractive to those researching within "noisy" realms, where the signal-to-noise ratio is low. Soft computing is normally accepted to include the three key areas of fuzzy logic, artificial neural networks, and probabilistic reasoning (which includes genetic algorithms, chaos theory, etc.).

The arena of investment trading is one such field where there is an abundance of noisy data. It is in this area that traditional computing typically gives way to soft computing, as the rigid conditions applied by traditional computing cannot be met. This is particularly evident where the same sets of input conditions may appear to invoke different outcomes, or there is an abundance of missing or poor-quality data.

Artificial neural networks (ANNs) are a particularly promising branch on the tree of soft computing, as they possess the ability to determine nonlinear relationships and are particularly adept at dealing with noisy data sets.

From an investment point of view, ANNs are particularly attractive, as they offer the possibility of achieving higher investment returns for two distinct reasons. First, with the advent of cheaper computing power, many mathematical techniques have come to be in common use, effectively minimizing any advantage they had introduced (Samuel & Malakkal, 1990). Second, in order to attempt to address the first issue, many techniques have become more complex. There is a real risk that the signal-to-noise ratio associated with such techniques may be becoming lower, particularly in the area of pattern recognition (Blakey, 2002).

Investment and financial trading is normally divided into two major disciplines: fundamental analysis and technical analysis. Papers concerned with applying ANNs to these two disciplines are reviewed.

BACKGROUND

There are a number of approaches within the literature that deal with applying ANN techniques to investment and trading. Although there appears to be no formal segmentation of these different approaches, this review classifies the literature into the topics proposed by Tan (2001), and augments these classifications with one more category, namely, hybrid. These categories of ANN, then, are as follows:

- Time series—Forecast future data points using historical data sets. Research reviewed in this area generally attempts to predict the future values of some time series. Possible time series include base time series data (e.g., closing prices) or time series derived from base data (e.g., indicators, which are frequently used in technical analysis).
- Pattern recognition and classification—These are attempts to classify observations into categories, generally by learning patterns in data. Research reviewed in this area involved the detection of patterns and segregation of base data into "winner" and "loser" categories as well as in financial distress and bankruptcy prediction.
- Optimization—This involves solving problems where patterns in the data are not known, often nonpolynomial (NP)-complete problems. Research reviewed in this area covered the optimal selection of parameters and determining the optimal point at which to enter transactions.
- Hybrid—This category was used to distinguish research that attempted to exploit the synergy effect by combining more than one of the above styles.

There appears to be a wide acceptance of the benefit of the synergy effect, whereby the whole is seen as being greater than the sum of the individual parts. This can be easily seen by inspecting the following reserach, which clearly shows that Hybrid techniques account for about one third of the research reviewed.

Further, the bias in this style of research toward technical analysis techniques is also evident from the research, with one third of the research pursuing the area of pattern recognition and classification. Technical analysis particularly lends itself to this style of research, as a large focus of technical analysis concerns the detection of patterns in data and the examination of the behavior of market participants when these patterns are manifest.

MAIN THRUST OF THE CHAPTER

This section briefly considers the characteristics of each of the four main categories described above. The selected papers were chosen as they are representative of current research directions, represent an important change in direction for this style of research, or represent a novel approach.

Research into Time Series Prediction

The area of time series predictions is normally focused on attempting to predict the future values of a time series in one of two primary ways:

- Predicting future values of a series from the past values of that same series
- Predicting future values of a series using data from different series

Typically, current research in this area focuses on predicting returns, or some variable thought to correlate with returns (e.g., earnings). Some researchers focus on attempting to predict future direction of a series (e.g., increasing from last known value, decreasing from last known value, no change). Research of this nature is essentially a classification problem and is discussed in the relevant section.

The following papers were selected and reviewed as they are representative of the current research in time series prediction. The authors of those papers are Chan and Foo (1995), Quah and Srinivasan (2000), Yao and Poh (1995), Hobbs and Bourbakis (1995), Austin et al. (1997), and Falas et al. (1994). The papers reviewed consider both fundamental and technical data. For example, Falas et al. (1994) used ANNs to attempt to predict future earnings based on reported accounting variables. They found no significant benefit using ANNs compared to the logit model and concluded that the accounting variables chosen were not appropriate earnings predictors. This conclusion represents one of the major problems encountered when working with ANNs, namely, their nonexistent explanatory capability. It is common to find conclusions of this type when reviewing ANN research, with noncorrelation often being reported as wrongly chosen input variables. Quah and Srinivasan (2000) used mainly accounting variables to predict excess returns (with limited success). Chan and Foo (1995) used ANNs to predict future time series values of stock prices and used these "future" values to compute a variety of technical indicators. The ANN produced showed particularly promising results. The authors concluded that the networks' ability to predict allows the trader to enter a trade a day or two before it is signalled by regular technical indicators, and this accounts for the substantially increased profit potential of the network.

In many ways, these two primary prediction methodologies relate to technical analysis strategies. For example, the use (and projection) of a moving average over a series of stock prices could be regarded as predicting future values of a series (the moving average) from past values of the same series. Indicators in technical analysis are often composed of a number of constituent data items, like price, volume, open-interest, etc. These indicators are commonly used to give indications of the future direction of price.

Research into Pattern Recognition and Classification

Pattern recognition techniques and classification techniques have been grouped together, as their goal is normally not to predict future values of a time series, but to predict future direction of a time series. For example, the primary goal of a chartist (a style of technical analyst) is to attempt to predict trend turning points by studying chart price action, looking for certain patterns. Chartists have noticed that these patterns tend to recur and are reasonably reliable indicators of the future direction of price trends. There are a number of these chart patterns, and different analysts attach different weightings to the predictive power of any given pattern. Also, these patterns normally need to be confirmed by values from another time series (such as volume) to be considered reliable. For more detail on this area, the reader is encouraged to refer to Pring (1999). Nonpattern matching techniques that also attempt to predict the future direction of a time series are also classification problems. Often, in addition to predicting the future direction of a time series, classification research attempts to classify stocks into two main groups, namely, "winners" and "losers," as in bankruptcy and financial distress predictions.

The following papers, listed by author, were selected and reviewed as they are representative of the current research in pattern recognition and classification: Tan (1996), Tan and Dihardjo (2001), Mizuno et al. (1998), Baba