Trade with Intelligence…

A TradingSolutions Primer: An overview of technical analysis using TradingSolutions' neural network technology.

NeuroDimension, Inc.
1800 N. Main Street, Suite D-4
Gainesville, FL 32609
1-800-ND-IDEAS
Outside US: 352-377-5144
Fax: 352-377-9009
www.tradingsolutions.com
info@tradingsolutions.com

A TradingSolutions Primer

TradingSolutions is a comprehensive financial analysis software package that helps you make better trading decisions by combining traditional technical analysis with state-of-the-art neural network technologies. It has the ability to learn patterns from historical data, allowing you to create highly accurate systems that inform you when to enter and exit positions.

This document is only intended to be a brief overview of TradingSolutions. A more complete source of information can be found at www.tradingsolutions.com. In addition, a free evaluation copy of TradingSolutions can be downloaded from this site. The evaluation version gives you full access to all of the software's capabilities using the included sample data. It also includes animated demonstrations, step-by-step tutorials, and daily trading signals from the new Solution Service.

Introduction to Technical Analysis

Technical Analysis involves using past stock prices, volume, and other related data to forecast future price movements. There are three basic premises that a technical analyst believes in. First, technical analysts believe that the price of a stock is driven by supply and demand. Stocks are not always worth the price that they are selling for. They often trade higher or lower based upon a large demand or a lack thereof. This ties into the second belief of the technical analyst: stocks move in trends that usually last for a detectable period of time. In other words, price movements are not simply random variations. The final belief of a technical analyst is that these detectable trends often repeat themselves. By detecting a repeating pattern in the early stages, a technical analyst is able to profit from the stock price movement if it behaves in the same manner that it did in the past.

There are two main types of technical analysis. The first type is based upon the recognition and interpretation of chart patterns. This type of technical analysis is more of an art than a science and can be very subjective since it relies on the technical analysts' judgment for determining whether a pattern exists or not. Common techniques utilized in this type of technical analysis involve drawing trendlines on the chart, interpreting a Japanese candlestick chart, and using other line studies such as Fibonacci Arcs, Gann Fans, etc. One major problem with this type of technical analysis is that it does not easily lend itself to historical backtesting. Because pattern recognition is subjective and since each pattern must be recognized manually by the analyst, historical backtesting is almost prohibitively laborious. See figure 1, below, for an example of how a trendline might be drawn on a recognized trend.

Figure 1: *Demonstration of a Trendline. Notice the distinct upward trend in the CSCO chart from November 5th to December 13th. This type of trend can be exploited by purchasing the stock after the trend has lasted for so many days and selling the stock once the trend appears to be broken.*

The second type of technical analysis uses mechanical indicators based on price, volume and other data to predict future price movements and/or to determine when to buy or sell. The data used to create an indicator can be technical (or even date-based fundamental) data taken from the stock being analyzed, a related stock or stocks, or from various general market data sources. This data is then applied to a

mathematical formula to generate the indicator. Some indicators can be used directly to make trading decisions. Other indicators must first be processed by a series of rules to generate buy/sell signals. The underlying formulas for many indicators are based upon moving averages, oscillators, etc. and use price data, volume data, and market data such as breadth as the inputs. Because this type of technical analysis is mechanical in nature and uses mathematical formulas, historical backtesting is possible. This allows the technical analyst to gain confidence in the indicators and any resulting signals before actually trading. See figure 2, below, for an example of an indicator.

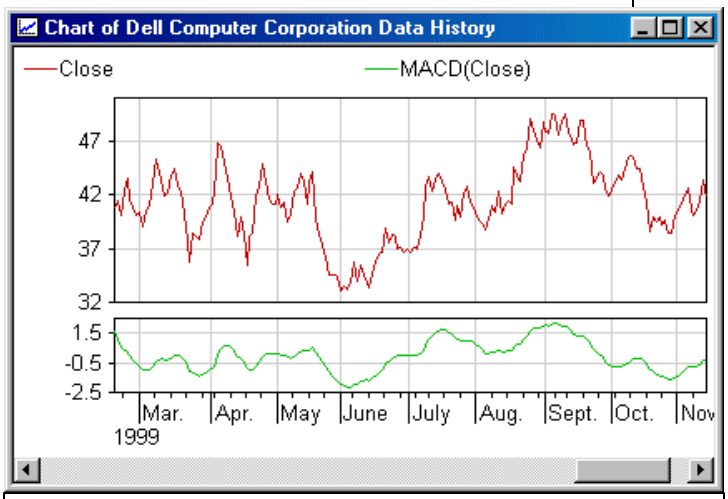

Figure 2: *Demonstration of the classic Moving Average Convergence/Divergence (MACD) indicator. When the MACD is increasing, this indicates that prices are trending higher, whereas a decreasing MACD indicates that prices are trending lower. The MACD is typically traded at the crossover points of its 9-day exponential moving average (not shown).*

TradingSolutions can be used for both analyzing charts and implementing mechanical trading systems. What makes TradingSolutions so much more powerful than other technical analysis software programs is its ability to create mechanical trading systems based on an advanced form of artificial intelligence known as neural networks.

Introduction to Neural Networks

A neural network is a powerful data modeling tool that is able to capture and represent complex input/output relationships. The motivation for the development of neural network technology stemmed from the desire to develop an artificial system that could perform "intelligent" tasks similar to those performed by the human brain. Neural networks resemble the human brain in the following two ways:

1. A neural network acquires knowledge through learning.
2. A neural network's knowledge is stored within inter-neuron connection strengths known as synaptic weights.

The true power and advantage of neural networks lies in their ability to represent both linear and non-linear relationships and in their ability to learn these relationships directly from the data being modeled. Traditional linear models are simply inadequate when it comes to modeling data that contains non-linear characteristics.

The most common neural network model is the multilayer perceptron (MLP). This type of neural network is known as a supervised network because it requires a desired output in order to learn. The neural network learns how to make the association between a set of inputs and a corresponding set of desired outputs using historical data. The goal is for the neural network to then be used to take in a new set of inputs and produce a useful output when the desired output is unknown. A graphical representation of an MLP is shown in Figure 3.

The MLP and many other neural networks learn using an algorithm called backpropagation. With backpropagation, the input data is repeatedly presented to the neural network. With each presentation the output of the neural network is compared to the desired output and an error is computed. This error is then fed back (backpropagated) to the neural network and used to adjust the weights such that the error decreases with each iteration and the neural model gets closer and closer to producing the desired output. This process is known as "training".

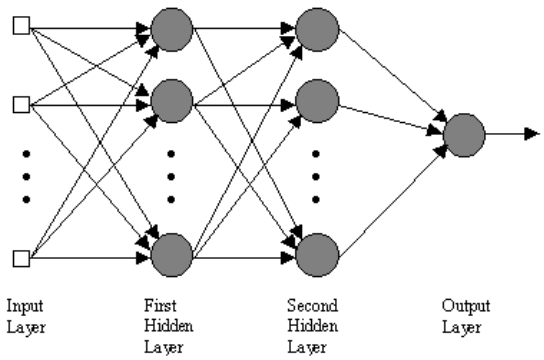

Figure 3: *Block diagram of a two-hidden-layer multilayer perceptron (MLP). The arrows represent interconnection weights and the circles are processing nodes.*

Using Neural Networks for Technical Analysis

Neural networks have recently become a very popular technology for technical analysis and they are at the core of TradingSolutions. They most closely fall within the bounds of the mechanical type of technical analysis described above. A neural network can be viewed as a very powerful and flexible "formula" whose output can be used directly for making trading decisions or to produce buy/sell signals. However, the neural network goes far beyond a typical indicator, because it has the ability to learn from the data itself. It is able to find relationships and patterns in the data that are often too complex to identify from reading a chart or to define with a set of mathematical rules.

The most common way neural networks are used for technical analysis is to create a model for predicting the future price of a financial instrument, given the current and previous prices and other technical and/or fundamental data. The predicted price could then be processed to produce a signal indicating when to buy or sell. The simplest system would be: buy if the predicted price is greater than the current price and sell otherwise.

Let's take a look at a simple example shown in Figure 4. Suppose you want to give a neural network today's closing price of a particular stock and the value of today's S&P 500 Index. What you would like the neural network to produce is a prediction for tomorrow's closing price of the stock. In order to do that you need to train the network using several years of data, giving it the stock closing price and index value for each day as the two inputs and the stock closing price of the next day as the desired output. You then would test the neural network on new data to verify that the prediction is good, and then apply a system to transform the predicted price into buy/sell signals.

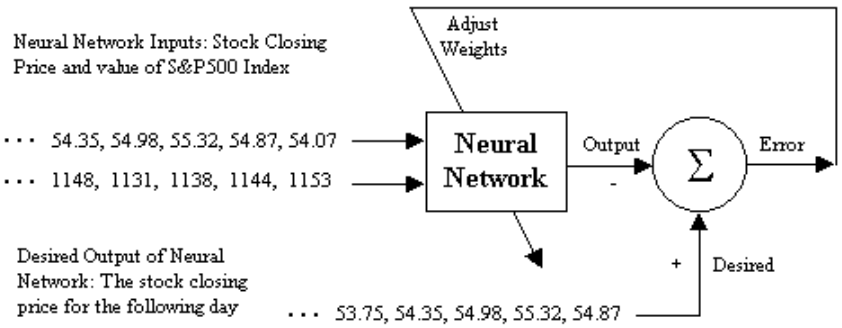

Figure 4: *Demonstration of a neural network learning to predict a future price of a stock. The stock price and S&P500 index data is repeatedly presented to the neural network. With each presentation, the error between the network output and the desired output (the next day's stock price) is computed and fed back to the neural network. The neural network uses this error to adjust its weights such that the error will be decreased. This sequence of events is usually repeated a few hundred times.*

A generally more effective approach supported by TradingSolutions is to train a neural network to produce the buy/sell signals directly instead of a future price. The desired output for the neural network training would be a set of numerical values representing "optimal" buys and sells, produced by looking into the future to determine the optimum time to buy and sell (see Figure 5). The neural network is then trained to produce these optimal buy/sell signals using only current and historical data as inputs. Once trained, the neural network signals from new data

can be used directly for testing and/or trading – there is no additional processing required.

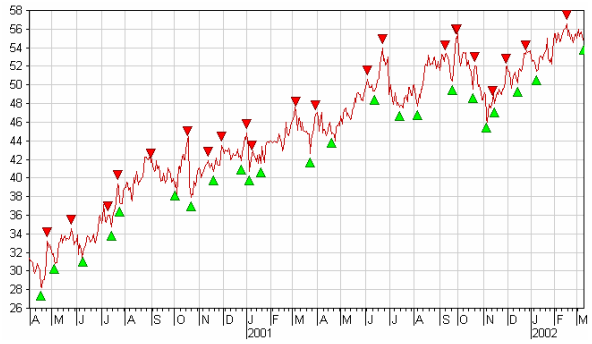

Figure 5: *"Optimal" trading signals for Baxter International (BAX). The green up arrows indicate long (buy) signals and the red down arrows indicate short (sell) signals. A numerical representation of these buys and sells can be used as the desired output for a neural network model.*

Using TradingSolutions to Create Profitable Trading Models

One of the key steps in developing a profitable model is to find a set of inputs that a neural network will be able to associate with the future price direction. There are over 200 built-in technical indicators that you can use, such as moving averages, stochastics and relative strength, as well as any custom indicators you create. TradingSolutions includes tools such as correlation analysis and genetic optimization to help you find the best data to use as inputs into the neural network model.

Figure 6 illustrates a sample TradingSolutions model for Baxter International (BAX) that was created back in March of 2001. You can see from Figure 6 that it has performed remarkably well ever since: 10 trades, 90% correct, 80% return vs. a buy-and-hold return of 19%*.

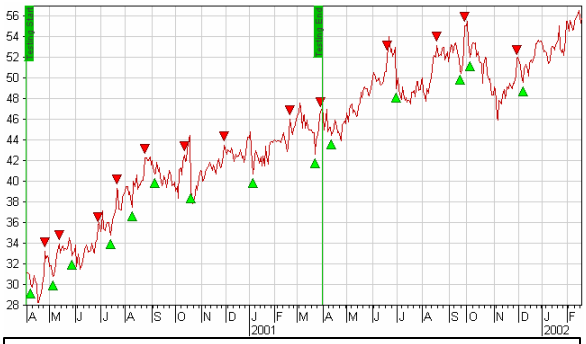

Figure 6: *Out-of-sample trading signals of a TradingSolutions neural network model for Baxter International. Notice this graph is not identical to the "optimal" signals shown in Figure 5, but it is close enough to be very profitable.*

There are only two inputs into this neural network model: 1) the percent change in the closing price, and 2) the percent change in the highest price traded over the past 10 days. The network was trained with an optimal signal as the desired output using the 4 years of data preceding April of 2000. The network was then tested using the data from the next year and the results were analyzed (see Figure 7). The performance of the model during that year was exceptional, so the signals were traded with real money from that point forward.

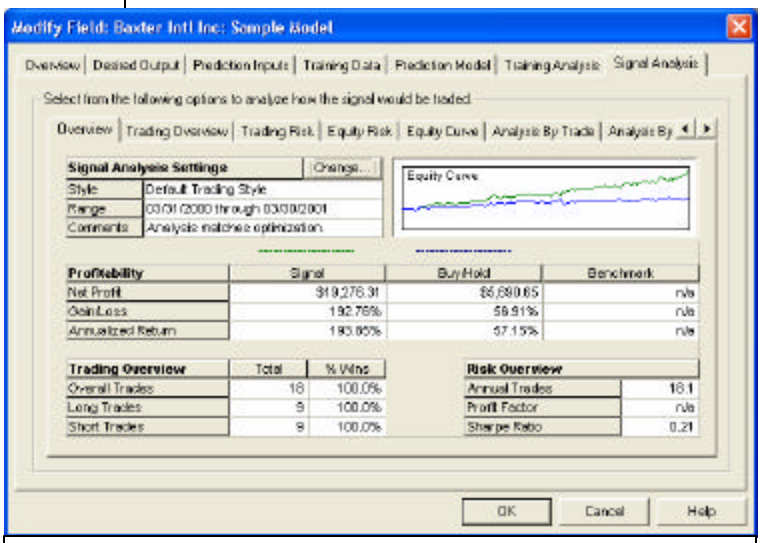

Figure 7: *Signal analysis statistics for the neural network model for Baxter International during the out-of-sample testing set.*

Taking Advantage of the Solution Service

It can take some time to learn how to create your own profitable neural network models. Until then you can take advantage of the free Solution Service provided to both the licensed and the evaluation users of TradingSolutions. The Solution Service will provide you with daily data, commentary and trading signals from several of NeuroDimension's top-performing neural network models, including the Baxter model shown above. Past performance of all models can be analyzed and displayed using the complete set of analysis and charting tools included with the software. The techniques demonstrated in these models can also be used as a starting point to develop your own models for other stock, futures, or options data.

Conclusion

Developing models or reading charts to predict the future behavior of a financial instrument is a very difficult problem. You will never be able to come up with a system that will make money with every trade. The advantage to using neural networks can be found in the complexity of the patterns they can identify. This includes patterns that cannot be seen in charts and patterns that could not be easily described through manually entered rules. TradingSolutions goes one step further in that it utilizes an optimal signal, which highlights the most profitable trades throughout the history of the stock in a way that makes that information easiest for the neural networks to analyze. This gives TradingSolutions users a significant advantage in creating successful trading systems.

Contact Information

NeuroDimension, Inc.
1800 N. Main Street, Suite D4
Gainesville, FL 32609
Toll-Free: 1-800-634-3327
Outside US: 352-377-5144
Fax: 352-377-9009
Email: info@tradingsolutions.com
Web: www.tradingsolutions.com

* Past performance does not guarantee future results. The financial markets are inherently risky. NeuroDimension is not recommending the purchase or sale of any security. The data provided is gathered from sources believed to be reliable, but no representation is made as to the accuracy of any data presented. Investors should use their own judgment in making investment decisions. NeuroDimension assumes no liability for any losses that may be incurred when using its software for trading.