

SUITS US, SIR

SIRTRADE INTERNATIONAL HAS COME UP WITH A
UNIQUE, TRADER-FRIENDLY WAY OF COMBINING
THE POWER OF NEURAL NETS WITH THE
SOPHISTICATION OF FUZZY LOGIC.

ANDREW WEBB REPORTS

For the trader hunting for a new trading methodology, artificial intelligence techniques have a particular allure. The concept of your artificial expert tirelessly perusing a great mound of data, identifying subtle relationships that would defeat the perception and stamina of a human equivalent is undeniably attractive. Unfortunately, though the situation has much improved in recent years, the snag with these techniques has been that many of the available packages were rather less than trader-friendly. Hours of pre-processing data and grappling with a user-hostile interface soon exhausted the limited attention span of the trader, and saw the concept (and software) swiftly consigned to the shelf or bin.

Pierre Orphelin, of SirTrade International in Paris, appears to have taken these problems very much to heart when designing the company's SirTrade neurofuzzy trading application (a combination of neural net and fuzzy logic) and its accompanying Assistant for Expert Traders. The two packages interact to allow the trader to develop neurofuzzy trading systems that combine the power of neural nets with the sophistication of fuzzy logic, and can be used (and even retrained) in real time in Omega's TradeStation. The only significant requirements on the user's part are a good understanding of technical analysis and more than a passing acquaintance with TradeStation's Easy Language and Power Editor.

In view of the significant amount of processing power that the Assistant's training routines require, Orphelin recommends that the two programs be ideally run on separate machines to prevent interference with real-time data flow via the comports. However, I encountered no difficulties running both applications on a Pentium Pro 200 with 128MB of RAM.

The first step is to produce some suitable data that can be used in 'training' the neurofuzzy expert trading system so that it can produce viable trading results in real time. This is accomplished in TradeStation using a template indicator supplied by SirTrade (that also includes a number of popular technical analysis indicators), and printing the output to an ASCII file. By modifying the Easy Language code you can include any selection of technical

indicators that you feel are of predictive value. Apart from the numerical output of the selected indicators for each period, the ASCII data also contains a target value that would trigger a trading signal for the indicator and the difference between the closing price of each time interval and its predecessor.

Once the ASCII file has been produced, one switches to the separate Assistant program for the training routine. When starting the Assistant, the first window displayed is the journal file, which records the results of all activities undertaken during operation. As a tool for debugging any training problems it is invaluable. Selecting 'File: New' from the menu bar and 'Fuzzy Description Files' from the ensuing pop-up menu opens the Command window, from which the ASCII file is selected and loaded. The window also allows the user to specify a wide range of training parameters. These include some very 'real world' applications, such as the ability to generate additional synthetic price bars and reinforce training under difficult (but vital) market conditions, such as the beginning and end of a trend. You can also specify how much of the available data is to be used for training and how much for out-of-sample testing – the default is a fifty-fifty split.

Having loaded the data file, the next step is to define the fuzzy description settings for the data that will be used during training. This consists of defining the min and max numerical values for each indicator (the software fills these in automatically after scanning the file, but you can change them if you wish), and specifying the desired number of fuzzy sets. To ensure effective training, this last should be set to at least the minimum number of different 'states' that the indicator can be in. Orphelin uses the RSI (Relative Strength Index) as an example that has three states – overbought, neutral, oversold.

After completing this stage, two new windows appear on screen – a Target Performance Summary and a Performance Summary. The former contains details of the hypothetical best results that could be achieved on the data set, while the latter will display performance results for training and test runs. The detailed performance analysis available here goes way beyond any other trading system analysis I have seen to date.

Having specified such para-



Get training!

meters as the number of iterations the software should make through the data and the training method (both back-propagation and stochastic are available), hitting the 'Train' command button sets the ball rolling. There must be some well-optimised code underlying the Assistant, as even training large, complex data sets with a large number of iterations selected didn't seem to take long (a great deal faster than some neural networks I have tested!). After training is complete, an equity curve is displayed, complete with buy/sell signals and price bars.

However, probably the most potent command in SirTrade is 'Explore', which examines multiple fuzzy sets to find the best performer for a given set of indicators. You realise just how much time this is saving as you watch the Journal window's scroll bar go berserk and multiple equity curves and performance reports flash before your eyes! Having selected the top performing FIS (Fuzzy Inference System), it can be further refined with additional training sessions or saved as an FZB file for use in TradeStation.

Having saved the file, TradeStation's Power Editor is again used with another template file provided with SirTrade, to complete the new trading system. At its simplest, this consists of little more than specifying the path to the FZB file. Load the trading system to a TradeStation real-time chart and away you go!

However, SirTrade offers you a lot more than just this. Using templates provided, you can specify a trading system that references two or more FZB files and Fuzzy Inference Systems. You could, for example, specify and train FISs for trending and

non-trending conditions, then specify a trigger switch that will automatically flip the trading system between them. If you want to be really smart, train multiple FISs and trigger trading signals based on their 'voting' consensus.

All this in itself is impressive enough, but Orphelin pushes the technology to the limit by also allowing real-time retraining of FISs from within TradeStation, so your system's expertise can be based on the most recent market price action. You can specify how many price bars elapse before retraining takes place (the default is five). Alternatively, you can specify conditions that will initiate retraining. The example provided with SirTrade uses the ADX indicator falling below a value of 15, which implies a non-trending market that will need more self adaptation than a trending one.

In brief, this is the most impressive application of artificial intelligence techniques to trading that I have seen to date. The power and flexibility are impressive, with pretty much every angle covered. Orphelin has clearly focused on developing a product that will get you quickly up and running with a viable trading system – not bogged down in side issues and data pre-processing. Though the performance results I achieved over the test period were anecdotal rather than rigorous, they were nevertheless extremely impressive. The methods contained in SirTrade are clearly viable. My only real criticisms are of the manual and help files, which are in the most exquisitely tortured English! Unravelling the ruptured syntax does make learning the ropes more long-winded than necessary. That aside, this is an extremely well-planned and effective product. ♦

BYTES

CREDIT (NUMBER) CRUNCHING

In a recent survey by the International Securities Market Association (see DELTA, p21), credit derivative market participants complained that a lack of available comprehensive data stops them using credit risk models. The signs are that the complaint is to be addressed.

Standard & Poor's has joined forces with Summit Systems to integrate its Credit Pro data tool with Summit's Credit Risk Engine, and Axiom's RiskMonitor. This will enable users to refine portfolio credit risk exposure calculations "dynamically", by establishing the parameters that most closely reflect securities held within a portfolio.

The alliances will set minds racing, not only at the other credit agencies, but also primarily amongst software vendors. For any firm lucky enough, a link with a major rating agency will provide kudos to a credit risk application.

Moody's offers its standalone Moody's Credit Risk Calculator and its Default Risk Service, a consultancy practice on default risk. The Calculator is a software and database package for assessing default and rating migration data, which generates tables of default rates via a consultative service. But the firm has no comment as to whether it is working with other parties on this type of alliance.

At the same time, Infinity has unveiled version 98.1 of its enterprise-wide risk management system, Panorama. Credit risk is the main area of focus for upgrades to the product – it now allows real-time Monte Carlo credit risk future exposure, and VaR calculations.

Counterparty selectivity is among the most topical of issues – inappropriately high risk loans having been blamed for the downfall of Peregrine, for example. Panorama's solution is

credit exposure profiling, which allows users to graph the potential future exposure for a selected counterparty at each term.

Users might like to match Panorama against Infinity competitor Cats' Carma to find the right product. With this in mind, Cats has benchmarked Carma's speed. The system works at 1.7 million revaluations a minute for 230 open positions, which approximates to near real-time enterprise-wide credit risk analysis.

BARRA ACQUIRES REDPOINT

Barra has acquired enterprise-wide risk management specialists Redpoint Software. The firm paid \$5.5 million immediately, with a possible \$12.5m of performance-based future payments.

The acquisition, which will see Barra accept certain liabilities of Redpoint, provides Barra with the benefit of access to Redpoint's enterprise-wide risk system, TotalRisk, and FinWorks toolkit. The benefit for Redpoint, according to Redpoint director of marketing Patricia Nezil, is access to Barra's worldwide presence.

Nezil comments that the products of the newly merged company will stay independent. The individual companies will also be able to ally with other data and analytics providers, according to Nezil.

Redpoint's previous agreement with Infinity, the aim of which was to utilise the strengths of the two company's toolkits, was scrapped. Nezil says the two toolkits involved in the alliance were competing products, and neither company saw the need to go forward. ♦

Keiron Greenhalgh