 By *Tapen Sinha*

Artificial Neural Networks

Biological neural networks are all-pervasive. Organisms as rudimentary as slugs have them; so do humans. All the slug has, though, is a simple set of instructions "programmed" into its nervous system. The central unit of the system is a neuron, each of which is connected to many other neurons via synapses. This network of neurons is what's termed a "biological neural network."

At the other extreme from the slug's humble nervous system is the human brain: It has 100 billion neurons, with roughly 100 trillion synapses. The involuntary functions of the human body (regulation of body temperature and heart rate, for example) are all controlled by the biological neural networks. Higher functions as well—our ability to see objects, understand speech, and think thoughts—are part of the biological neural network, too.

In some sense, an artificial neural network (ANN) mimics its biological counterpart. Still, the differences between the two are substantial. For one thing, current computer elements are a million times faster than the biological neural switch time. On the other hand, though, there's a thousand times greater connectivity between the neurons in the human brain than in today's supercomputers. Second, the processing of information in the human nervous system is carried on in several layers of networks at the same time. In computer jargon, the nervous system achieves massive parallel processing; in computers, we're still a long way off from achieving truly massive parallel processing.

Nonetheless, the performance of ANNs in various areas has been substantial. Here, I'll explore the applications of ANNs in actuarial work, in nontechnical terms. (Readers interested in a more technical exposition should refer to my article in the January 1994 issue of *Digital Doings*, published by the Society of Actuaries.)

At its most basic level, an ANN consists of a connection between a set of inputs and a set of outputs, via a set of weights. The output is obtained by applying the set of weights to the inputs and applying some (nonlinear) transformation. Typically, the inputs and outputs have an intermediate link: the hidden layers. Most often, there is only one hidden layer and only one output. For most uses, a single layer of hidden units is enough to produce a good model.

ANN is used to recognize patterns in the data and to make predictions. At its simplest level, it serves as an alternative to statistical methods. In the Gulf War, ANN was used extensively by computers on board aircraft and missiles for identifying targets. Medicine makes extensive use of ANN—diagnosing cervical cancer, for instance, from the results of medical examinations has successfully used methods of ANN.

Fidelity Investments pick 200 stocks out of a universe of 2,000 using an ANN-based program. Their selection beat the Standard & Poor 500 by 2% to 7%.

ANN in Stock Market Predictions

Roughly speaking, if averages matter, but volatility doesn't, then usual statistical methods are good enough. However, if volatility of the set of variables under scrutiny *does* matter, then ANN provides a better alternative—and researchers in finance have recognized that variability does matter. In standard statistics, we can deal with only one problem at a time. Thus, ANN has an advantage, since it takes a more holistic approach in finding all patterns that manifest from nonlinear connections *simultaneously*.

Forecasting in the stock market is a growing area of applications of ANN. Some academics argue that any ANN program (for predicting a market index or a specific company's stock price) has to overcome objections from modern portfolio theory, which implies market efficiency. The immediate implication of market efficiency is that it's impossible to make a profit by using past information from time series alone.

However, in the stock market, we observe traders paying immense attention to "technical" information. In fact, most of the newspaper reports we see every day about the stock and other asset markets is technical information (in phrases like "the market is overvalued," or "the dollar crashes through the support level of 100 yen").

If we carefully document the "evidence" in favor of stock market efficiency, we find (1) the evidence is mainly about not rejecting the null hypothesis about market efficiency, and (2) the tests are all joint tests of linearity and normality of the model. In other words, there may be nonlinear (and complex) models that might be able to predict the future values of share prices from past prices.

Early in the resurgence of ANN research, *physicists* from Los Alamos research laboratory were speculating about applying ANN to analyzing stocks! In the first serious application, Halbert White of the University of California, San Diego, used daily price data for IBM stock, and his results were inconclusive. But later studies showed that *weekly* data provide a more meaningful set of predictions; many commercial programs now offer stock market prediction as part of their ANN packages.

How successful have ANN-based programs been? Most traders are reluctant to talk about it publicly. However, Fidelity Investments picks 200 stocks out of a universe of 2,000 using an ANN-based program. Then, they move on to human experts, who narrow the field to fifty, in which they invest $2.6 billion (reported in *The Economist*, October 9, 1993). Their selection beat the Standard & Poor 500 by 2% to 7%, per quarter, over 3 years.

Fijitsu, in Japan, has experimented with applications of ANN in the stock market. Here's what they've found so far: Their ANN program beat the market index in Japan, to some extent, during a boom, and much better than the index during a bear market. This precisely replicated the pattern for human analysts. So in this sense, over a period of several years, the ANN programs are mimicking human experts quite nicely.

ANN in Detecting Health Insurance Fraud

Australia has a universal basic health insurance plan (called Medicare) that's funded by the federal government through a specific tax (the "Medicare levy") and by tax revenue from the general account. Citizens and permanent residents of Australia are all issued Medicare cards, which serve as proof of eligibility for benefits. Most medical practitioners send their bills directly to the federal government, and get reimbursed for the services rendered.

Some medical practitioners demand payment at the time of the visit to the doctor. In that case, the patient can claim a predetermined amount by providing documentary evidence at one of the Medicare offices, which are located in shopping centers around Australia. For claims under $500, cash payments are made on the spot. A subsection of the population (retirees, for instance, who get age pensions, and the unemployed) get additional benefits in the form of reimbursement for prescribed medication, etc.

There are 36,000 general practitioners (plus a few thousand specialists) in Australia. Each year, the total number of claims processed by Medicare exceeds 200 million. Some anecdotal evidence suggested that some doctors were inflating the value of the services they'd provided. So the federal government decided to investigate.

ANN has an advantage, since it takes a more holistic approach in finding all patterns that manifest from nonlinear connections simultaneously.

At the outset, it wasn't clear how many doctors were involved. It soon became clear that in fact only a small fraction of all doctors were inflating bills—but not at the same rate. So the next step was to create a profile of a "normal" bill. And since the various kinds of services have different fees, categories of services had to be set up. The next step was the formation of a "cluster" of normal bills. Once that was done, it was possible to spot the bills that don't belong to the normal clusters.

For each category, a specific kind of ANN, termed an "associative memory network" (or, Kohonen network), was created to form the clusters. The model created an "ideal" path of a transaction with no inflated figures. Then, the remaining task was to catch the doctors who'd inflated their figures.

ANN identified approximately 500 outliers. Some of them were further off the norm than others; these doctors were identified and then audited. Many had no acceptable explanation for their exorbitant charges!

The statistically minded reader will immediately realize that this feat of detection could have been accomplished using a statistical method of cluster analysis. In principle, it's a lot like the anthropological problem of classifying a newly found set of bones to classes of bones people already identified as anthropometrically similar. Usually, these kinds of problems involve only a small body (no pun intended) of data. But that wasn't the case here. The time required to analyze such a large amount of data using cluster analysis would have been considerable. Just the processing time required to generate a list of outliers among 200 million records, using statistical techniques, would be enormous.

ANN made the process simpler, easier, cheaper, and faster.

Is ANN a Fad?

In the past, two major problems have plagued ANN. First, it's an interdisciplinary subject. So when the PET scan designers, applying the analogy of contraction of muscle tissue to traders in the stock market to predict market behavior, they might miss something obvious to researchers in finance, unless they'd had some prior training in the subject. In consequence, ANN has suffered some spectacular failures.

Second, most of the applications of ANN are "black box" applications: ANN works in diverse fields, but the underlying mechanism as to *why* it works isn't well understood. That makes many analysts very nervous!

But one sign of acceptability of ANN comes from academic researchers in finance, who are otherwise dismissive of claims of market inefficiencies. In 1993, there was a major conference of researchers focusing exclusively on the applications of ANN in finance at the London Business School. This year, a similar gathering is taking place at the California Institute of Technology.

In a very fundamental sense, ANN is a product of recent advances in the speed of computers. If computers ran at the same speed as 20 years ago, implementation of ANN in "real time" would be impractical. However, even today's computers have trouble implementing ANN because all computers process information serially. *Biological* neural networks, on the other hand, process information in parallel. Thus, we won't be able to realize the full potential of ANN until the massively parallel processors (the so-called sixth generation computers) come into existence. □

***Tapen Sinha, Ph.D.**, is associate professor of finance, Bond University, Queensland, Australia.*

PUZZLES

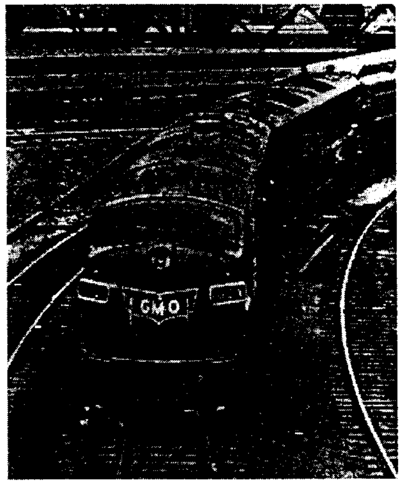

1. (A) Prove that, with one exception, the product of two positive whole numbers, each greater than one, exceeds their sum. (B) Now, prove you were wrong, that **6 x 6 = 12**, showing a dozen steps, more or less, along the way, in this letter multiplication puzzle from *The Actuary*, a publication of the British Institute of Actuaries.

```
      S I X
      S I X
    -------
    _ _ I _
  _ _ _ _
_ _ A S
-----------
A D O Z E N
```

2. Three new companies together leased a branch line from a railroad at a fixed monthly rental. They agreed among themselves to share the monthly payment in proportion to the number of cars they respectively shipped each month. In the second and third month, each company shipped two more cars than in the month before. Calculations at the end of three months showed that the first company's share for the first month was **$120**; the second company for the second month, **$135**; and the third for the third month, which shipped five cars in the first month, was **$108**. How many cars did each ship the first month?

***Please send all answers to:** Ralph E. Edwards, Editor, 3024 N. Calvert St., Apt. A-1, Baltimore, MD 21218. The names of those who send correct solutions will be acknowledged in a following issue of* Contingencies. *We do not have space to print solvers' solutions.*

Answers to September/October Puzzles:

1. *The question was:* A dealer with a balance scale has **100 weights** to use with it. He wishes to retain as few of them as possible but be able to weigh any integral amount up to **100 units** (pounds, ounces or whatever). He also wishes the heaviest one retained to be as light as possible. Which weights should he retain (A) using only one pan for the weights; (B) using both pans? Modify these results (C) and (D) if he says he never sells just one unit, and a one-unit weight is so small he prefers not to have one.

The answer is:

(A) **1, 2, 4, 8, 16, 32** and **37**
(B) **1, 3, 9, 27** and **60**
(C) **2, 3, 4, 8, 16, 32** and **37**
(D) **3, 5, 9, 27** and **64**

2. *The question was:* A farmer with two belligerent goats must devise a way for them to graze an entire rectangular meadow (**10 x 17½ yards**) without being able to fight. Two poles are **5 yards** from the sides and end of the field and **7½** yards apart. How can the two goats be tethered to the two poles with rope and rings so they are separated but have access to all the grass? How much rope is needed for the job?

The answer is: The rope is **14½** yards long and passed through both rings. (This assumes that **7** is the square root of **50**).

***Solvers (Previous Issues):** G. Crosby, M. Fowler, A. Goldberg, S. Gut, R. Hupf, O. Reed, D. Skurnick, J. Tupper, J. Ward, G. Watson, R. Wilton.*

***Solvers (September/October):** A. Goldberg, J. Kiel, F. Knorr, R.C. Martin, E. Torrance, J. Tritz, J. Ward.*