

IDEAS

Computers you can train work somewhat like a brain

Scientists are developing neural networks that enable computers to 'learn' functions

By **Simson L. Garfinkel**
Special to The Christian Science Monitor

THE mind-vs.-brain controversy is coming to a head in the world of artificial intelligence and computer science.

While researchers admit they're far away from a machine that can think - estimates range from 20 to 50 years or even never - computers already developed replicate many tasks, such as learning, once the sole province of biological brains.

"In AI there are two goals," says Bernardo Huberman, a scientist at Xerox's Palo Alto Research Center. "One is to build programs or to construct machines that behave or mimic the behavior of intelligent beings." The other is to use the creations to give better insight into the workings of human intelligence.

The debate is about the best way to proceed: by building electronic minds or electronic brains. Modeling the mind by developing computer representations of facts, concepts, and goals is a technique championed by a group of researchers called "symbol processors."

Another approach advocated by scientists called "connectionists" is to model the brain itself, building networks of electronic neurons and training them.

For 30 years, the symbol processors have consistently beaten out the connectionists in the battle for research funds, credibility, and graduate students. But recently, connectionists have found uses for their mathematical model of the brain - called a "neural network" - in applications as diverse as automated manufacturing, credit ratings, computer speech and vision, and the control of weapons systems.

More than 600 papers were submitted in September at the first meeting of the International Neural Network Society, an organization founded in March 1987 that already has 3,500 members from 49 of the United States and 38 other countries.

"Basically, it represents a new intellectual coalition," with members from psychology and neuro-science, mathematics, computer science, engineering, and business, says Stephen Grossberg, the society's founder.

Neural networks can quickly draw conclusions from large amounts of information. They can tolerate inaccuracies in the data they are given, and to some extent they can learn by trial-and-error or example.


Evaluating loan applications

Small networks of a few hundred neurons can be simulated easily on a personal computer. Adaptive Decision Systems, a small company in Andover, Mass., has developed a simulation that uses the neural network model to evaluate consumer loan applications. The network is first "trained" by showing it 5,000 credit applications - and identifying which loans were defaulted on. It can then look at new applications and predict the likelihood of a loan's being repaid.

The program works by identifying which characteristics of the applications are associated with good loans and which are associated with bad ones. The system makes mistakes - it sometimes approves loans that are later defaulted, or rejects loans that were actually a good risk. This is one of the problems with neural networks in general: They are never guaranteed to give the correct answer. But the program makes mistakes "much less often than people do, who rely on judgment," says the company's president, Murray Smith.

Last year, the Defense Advanced Research Projects Agency funded a \$700,000 initial study at the Massachusetts Institute of Technology's Lincoln Laboratory. Using the conclusion of that study, DARPA decided to fund a 17-month, \$33 million effort to further study neural networks, strengthen theory, and develop hardware.

The military might use neural networks for classifying images



Using an electronic 'brain' and camera 'eyes,' a robot puts candy in a box held by Michael Kuperstein

on a radar scope, says Alfred B.Gschwendtner, who headed the Lincoln Lab study. Today human operators classify such images, discriminating civilian aircraft from fighters, for example. Writing computer programs to do the same function has been very difficult, because people don't understand exactly what the operators do when they make their determinations.

Learning by example

In the future, says Dr. Gschwendtner (pronounced Schwidner), a neural network might be apprenticed to a human operator, watching the human's actions and slowly learning how to make the decisions. When it had seen enough examples, the network would begin to offer its own suggestions to the operator. These suggestions would become more accurate as time went on. Neural networks could also be built into missiles to guide them to their targets automatically.

"I think what people mean by 'smart weapons' is something like this," Gschwendtner says.

Outside the military, scientists have combined neural networks with television cameras and robot arms to create machines that can "learn" to move about and pick up objects in

much the same way a baby does - by repeatedly trying motions and watching the results.

Michael Kuperstein, a Brookline, Mass., inventor, has developed a system that will reach for and grab a ping-pong ball held in front of it. He holds the patent on the network that does the task, and his company, Neurogen, is using the technology to develop a robot for automatically placing chocolates in candy boxes. Martin Marietta, an aerospace engineering company, has developed a similar system, which can pick up skids from a conveyor belt using a neural-network-controlled forklift.

"It's almost an irresistible approach," says Thomas F. Knight, a professor of computer science at MIT. "You build a general-purpose machine, you put education in and get behavior out."

Neural networks are attractive because they hold the promise of being able to work on very large problems by solving different parts in parallel. There is also the promise that large computers constructed out of neural networks might be able to function even if parts of the network were destroyed, the same way people with small amounts of brain damage can often continue to lead normal lives.

The problem with neural networks, Dr. Knight says, is that they offer little insight in helping to understand the problems they are designed to solve. Unlike "expert systems" based on symbol processing, there is no way to "open up" a neural network and change the rules it uses to make decisions.

"Suppose the federal government just changed home loan interest so that it is no longer tax deductible," says Knight, when asked about the credit-predicting neural network. "How do I change my home loan system so that I tell it that the ground rules have changed?" The problem, he says, is that there is no single place within the network where the computer calculates the amount of money the mortgage buyer has available to spend.

Most neural network systems used today have a relatively small number of neurons - no more than a few hundred - and simulate them with special programs that run on conventional computers. Large-scale systems require networks with tens or hundreds of thousands of neurons, which will have to be constructed out of hardware designed specifically for that purpose. That hardware - now under development at Bell Laboratories, Lincoln Laboratory, and many other places - will package 50 to 200 neurons on a semiconductor chip, each of which will be a little analog computer capable simply of adding and multiplying a few voltages.

Neurogen's ping-pong-picking robot takes three seconds to make its decision on how to move the robot arm. With special neural network hardware, Dr. Kuperstein says, it could make a decision in the blink of an eye.

How neural networks work

Neural networks are modeled loosely on the human brain. A neural network consists of a large number of very simple electronic processors with a tangle of connections between them. Many researchers use the term "neurons" to describe these processors; others, who resist the brain analogy, prefer the term "nodes." The processors have a single function: to take their inputs, add them together, and produce a final result. The input of each processor is multiplied by a number called a "weight" before it is summed.

Most neural networks arrange the

processors in rows called stages. Data from the outside world enter the network through the first stage. The value of the second stage of the network is based on the results of the first stage. The number at the second stage can then be associated with some physical measurement, such as "position to move arm to."

Before it can be used, the neural network must be "trained," a process that involves finding which weights cause the network to give the desired answer when a certain input is presented to it.

For example, the Neurogen ping-pong robot is trained by placing a ping-pong

ball in the robot's grasp and instructing the arm to move to a random position (the arm is moved by a second computer). The network then correlates the image of the ball in each of the two video cameras with the arm's posture. This process is repeated 5,000 times, at the end of which the network is able to generate a desired posture for any position of the ball within the visual field.

For applications that require making decisions quickly on thousands or millions of pieces of information, neural networks offer possibilities unmatched by traditional computer systems. - S. L. G.

Science columnist Robert C. Cowen is on vacation.