THE SPANG ROBINSON REPORT ON

Artificial Intelligence

The Artificial Intelligence
Business Newsletter

September, 1987 Vol. 3, No. 9

Neural Networks: The New Frontier

The Spang Robinson Report tracks Artificial Intelligence technologies in both commercial and research environments. Each month, the report focuses on a key market or research issue furthering the commercialization of these technologies.

Neural networks are the latest challenge in simulating human cognitive processes on the computer. Artificial neural networks supposedly mimic the neural structure of brains, human or animal. There are certain leaders in Artificial Intelligence, particularly Marvin Minsky, who believe that not enough is known about the brain to accurately implement neural networks. From a scientific perspective, this is true. From a business perspective, this doesn't matter. The neural network approach is an alternative to traditional computing based heavily on well-known mathematics and can potentially solve problems that traditional computing cannot.

Already the market for neural network products is no more than $5 million from our own estimates, but there are nearly 100 research efforts going on worldwide. Products, applications and markets will begin to evolve.

In This Issue...

The Spang Robinson Report combines two previous AI newsletters: **The Artificial Intelligence Report** and **Target**. The Report is published monthly by Spang Robinson, P.O. Box 1432, Manchester, Massachusetts 01944, U.S.A. (617) 526-4820. Subscription rates are $295 in the U.S. and Canada; $345 elsewhere. Copyright:

Artificial Neural Networks

by Joseph Corrado

A new promise

Simulating an animal brain's neural structures is an idea that has finally found a market. Continuing demonstrations of artificial neural networks (NN's) have proven that enough theory is now mature enough to produce practical results. And more than merely practical, NN technology holds the promise of being able to solve several critical problems that have resisted von Neumann and conventional AI computing approaches for the past three decades.

Basic theory

In engineering terminology, a neural network is a highly parallel dynamic system with the topology of a directed graph. NN nodes are referred to as "processing elements", and the directed links (information channels) are called "interconnects". Each processing element receives multiple inputs and generates a single output signal that branches into multiple copies that are sent to other processing elements as input signals.

Differential equations

Each processing element's operation is determined by differential equations that define how the output signal develops in time as a function of the input signals. These equations are called the "transfer function" equations of the processing element. Other differential equations ("learning laws") can be used for modifying adjustable coefficients in the processing elements' main transfer function equations. Hence, a complete artificial NN can be described as a large system of coupled, ordinary differential equations.

Common paradigms

There are about a dozen types of neural network models or paradigms that have been defined so far. Each type is particularly well suited to a class of problems. Popular paradigms include the Kosko Fuzzy Cognitive Map (knowledge processing applications involving causal relationships among objects), and the Rumelhart/Williams Back propagation network (control applications such as speech synthesis).

Neural software

Despite early funding problems, neural network research persisted throughout the 1960s and '70s. In 1972, two physicists from Brown University formed an R&D partnership to create a computer model that could display some properties of the human learning process. This partnership evolved into Nestor Incorporated in 1975, and became the first publicly held company working exclusively in neurocomputing. With an innovative technology, solid products, a strong marketing orientation, savvy management, and $6 million in funding from the Reliance Group, this is a company on the move.

New hardware

Just as a LISP processor is required to optimize the processing of LISP code, a new type of computer is necessary for large-scale neural net computing. In contrast to a von Neumann computer architecture in which a central processor and memory manipulate data in a serial fashion, neural net processors or "neurocomputers" process data in parallel -- similar to an animal's nervous system in which each neuron is connected to perhaps a thousand others. With this type of architecture, pattern recognition is like the process of a liquid flowing immediately to the lowest points on an uneven surface -- organic, not mechanic.

Motorola

By 1980, important advances in neural network research had demonstrated that NN technology was not a dead end in the AI quest. About this same time, Robert Hecht-Nielsen, an engineer at Motorola, started a NN development program at Motorola that he took to TRW in 1983. That program culminated in the development of the first NN simulator--the TRW MARK III.

TRW

TRW's Mark III, with eight monoboard processors, offers a processing speed of about 500K interconnections per second (i/s). With 15 boards the figure becomes 1M i/s. With the Mark III as a fast prototyping tool, TRW hopes its customers will develop NN applications that can be implemented via TRW's VHSIC technology. This vertical integration strategy should serve TRW well in leveraging their technology assets. It also keeps them out of direct competition with the Hecht-Nielsen Neurocomputer Corporation (HNC).

Neural IBM

Departing from TRW's Rancho Carmel Artificial Intelligence Center in late 1986, Robert Hecht-Nielsen and Todd Gutschow formed HNC to pursue the full commercial development of neurocomputers. Their primary product is the ANZA neurocomputer -- a general purpose, board-level neurocomputer capable of implementing any neural network. HNC also offers standard "netware packages" which include different general-purpose and special-purpose network descriptions, and the AXON NN description language. Hoping to become the IBM of neurocomputing, HNC views NN as the technology that will make computers truly easy to use as well as powerful.

Market size

HNC estimates that NN will become a $100 million a year industry by 1990. By the year 2000, HNC believes that half of all computer robotics will be based on NN technology: a $250 to $300 million market for NN in robotics alone.

SAIC

Another player in the NN hardware market today is Science Applications International Corp (SAIC), a $600 million a year engineering design corporation that offers a very high-speed NN processor board (10M i/s doing Forward Propagation) called the Delta-1. They also market a workstation, the Sigma-1, that integrates the Delta-1 with a 386-based PC. Software bundled with these processors includes eight, ready-to-use NN programs controlled by a sophisticated user interface.

Market stages

Like many observers of the NN scene, HNC and SAIC expect the market to evolve in three primary stages. The first (current) phase focuses on the development of general purpose, digital neurocomputers able to implement all NN paradigms. In phase two (1989-90), application-specific neurocomputers based on true neural chips (custom analog IC's for high-volume applications) will be introduced. Custom analog chips for very high throughput systems will mark the advent of true neurocomputers -- not merely digital simulations of neural networks (the current state of the art). The next-generation neurocomputers are expected to enable real-time speech and image recognition, and very intelligent robotics.

The optical phase

During phase three (1992-97), optical neurocomputers offering billions of processing elements will begin to appear. These optical systems may quickly render many stage two machines obsolete because of their greatly-increased power and decreased cost. But even in stages two and three, HNC predicts that there will still exist a need for general purpose neurocomputers for continuing R&D and as delivery platforms for applications that may require more than one type of NN architecture.

New chips

Development of true analog NN chips is now underway at companies like Texas Instruments, Siemens, AT&T Bell Laboratories, and Synaptics Inc. At Bell Laboratories, scientists have already developed electronic neural networks (ENN's) composed of many simple components like resistors and amplifiers. AT&T's most complex chip to date has 256 electronic neurons composed of about 25K transistors and more than 100K resistors--all on a chip of silicon about 1/4" square. Another company to watch in the ENN chip area is Synaptics of San Jose, CA. Synaptics has attracted the attention of several venture capital firms who foresee immense paybacks in NN-based vision and speech-recognition systems.

Optical processing research

In the Adaptive Systems Group of Verac Inc. (San Diego, CA), Bart Kosko is doing basic software and hardware NN design research. Kosko believes that commercial NN optical processors will begin to appear in about five years. With this kind of timetable, systems based on VLSI neural chips will quickly be passed over (except for dedicated applications). With enormous data throughput rates, optical NN processors could prompt a computing revolution with unpredictable consequences.

Verac

At Verac, Kosko's program is dubbed ADAM (Air Defense Associative Memory). He notes that the only known programmable dynamic systems are neural networks. Hence they lend themselves to complex, dynamic, applications such as those typically encountered in military areas. Rest assured, the military has its sights set on advanced optical computing techniques.

In place of expert systems

Kosko claims that the most lucrative, current NN applications are large, real-time military programs where some of the funding -- for example in the Pilot's Associate project -- is being diverted from classic AI research. Kosko explains that all expert systems are basically search trees incapable of representing feedback and impossible to run in real time if the knowledge base is large. Moreover, he observes, all search trees are non-combinable. If the application is real-time and the knowledge domain is causal -- involving feedback -- large search trees (expert system approaches) simply cannot work, he says.

But no audit trail

Although the NN approach cannot provide an audit trail, Kosko notes that the issue becomes "who can predict most accurately -- not who can display a logic audit trail. Expert systems work best when they are small. For large-scale problems, a natural solution is hybrid systems of knowledge nets and expert systems."

Handling the unknown

John Voevodsky at Neuraltech Inc. (Portola Valley, CA), agrees. For the established AI tool vendors, does the marriage of expert systems with NN's make sense? Voevodsky declares emphatically, "Absolutely. What NN's allow you to do is handle the unknown without previous definitions of the unknown -- that's the bottom line." And what about the lack of an audit trail in NN-based systems? Says Voevodsky, "Data inputs can be ranked according to importance -- the way humans explain their decisions based on multiple factors." Neuraltech markets the Plato/Aristotle NN expert system development toolkit -- ***a tool that the established AI tool vendors would do well to investigate.***

Smarter expert systems

Casey Klimasauskas at NeuralWare (Sewickley, PA), also sees advantages in melding expert systems with NN systems. "There are some very powerful ways that expert systems and neural computing can be meshed to provide smarter expert systems," he says. ***"I think that research into integrating neural computing and expert systems is a very interesting strategy for a company to pursue."***

Wetware

For very large NN systems, biological neural matrices (maintained in a culture) connected with a mechanical interface is another future prospect for NN's. This "wetware" is already under investigation at places like AT&T Bell Laboratories and Martingale Research of

Allen, TX). Martingale currently has a contract with the Air Force (Wright Aeronautical Laboratories) for a project that involves research into how a biological neural network grown in a culture and an artificial neural network might be able to communicate with each other. Out of this research, Martingale is hoping to develop "neural prosthetics" -- devices for supplementing the abilities of sensory-deprived people.

Outlook

Verac's Kosko thinks that NN's are now butting heads with much classic AI work. "What is the main theorem of AI?" he asks rhetorically. "There is no theory, no testable hypothesis. At least, at its root, NN's are mathematical -- the basic principles of information processing stated in mathematical form." Kosko's advice to the AI tool and application vendors who want to break into NN's: ***"Review ordinary differential equations and matrix algebra."***

Potential problems in paradigms

Yet, the NN approach is not a panacea. NeuralWare's Klimasauskas believes the Achilles heel of NN technology may be difficulties in developing NN paradigms -- such as multi-layer learning algorithms -- for complex applications. He predicts: "If the industry cannot show that it is able to build better algorithms over the next four or five years, there is a chance that it might all go bust again, as it did in the late 60's."

New limits

In contrast to classic AI technology, artificial neural networks today are at the frontiers of a technology in which the limits of formal, sequential logic processing can be transcended in the simultaneity of a mere snail's next move. ***The implications -- and prospects of progress -- in neural network research are now so compelling that classical AI companies would be wise to prepare.***

(Other companies working in the area of neural networks include: AIware, Asahi Chemical (Japan), BDM, Bendix, DEC, General Dynamics, General Electric, Hughes Aircraft, IBM, Neural Systems (Canada), Neuronics, NHK Laboratories, Siemens, Revelations Research (Canada), Texas Instruments.)

Neural Net Report

For a detailed report on neural networks contact the Graeme Publishing Co. of Wilbraham, MA. They have published a report dealing primarily with neural network research activities which includes a description of commercial efforts. The report was written by Richard Miller.

Joseph Corrado is a writer in Palo Alto, CA. He consults on business and technical communications in AI and other fields of high technology.