Proc. of the 1991 IEEE
Int. Conf. on Tools for AI
San Jose, CA-Nov. 1991

PANEL: Are Neural Networks a Tool for AI?

Moderator: José G. Delgado-Frias, *SUNY-Binghamtom*

Panelist

Dan Hammerstron *(Adaptive Solutions, Inc.)*, Cris Koutsougeras *(Tulane University)*, Gerald G. Pechanek *(IBM)*, and Benjamin W. Wah *(University of Illinois, Urbana)*

In this panel the role of the artificial neural networks in the field of artificial intelligence will be discussed. Some other related topics such as artificial neural networks computational capabilities, the need for novel computer architectures, and neural network learning will be addressed as well. Each of the panelist have expressed his opinion on these issues. Brief statements by the panelists are provided below.

Dan Hammerstrom

I believe that neural networks are a part of AI and extend it in the direction of continuous computation (versus the traditional "discrete" methods). Traditional AI has had difficulty solving basic problems in speech and vision that involve the processing and reduction of massive quantities of loosely correlated data. Yet even basic animal nervous systems, which use switching devices are that are millions of times slower than today's transistors, solve these problems quickly and efficiently. It is clear, therefore, that the computational model used by these organisms is radically different from our traditional discrete computational techniques. Artificial neural networks attempt to capture some of those computational differences. They are only crude facsimiles of biological neural networks, but there are similarities.

And there are things that neural networks do not do as well as older methods, for example, the representation of complex, structured, "contextual", knowledge. This is where traditional AI has evolved a number of powerful techniques. Consequently, many researchers in the neural network field are interested in merging the structures and functions of these two approaches. The result of such merging would be processing systems with "hybrid vigor" that would leverage the best of both worlds.

Therefore, the neural network/connectionist computational model is complementary to, rather than competitive with, traditional Artificial Intelligent techniques -Artificial Neural Networks constitute an enhancing not a replacing technology.

And because of their rather different computational model, neural networks can benefit considerably from customized architectures. In fact because of their low precision and naturally parallel structures, neural networks lend themselves well to analog and hybrid analog/digital architectures. The ability of current VLSI to provide large numbers of simple processing elements (both analog and digital) cheaply allows for a quantum improvement in the cost-performance of neural network emulation.

Cris Koutsougeras

In addressing the question of whether neural networks are tools for AI one would have to approach the issue of defining intelligence. Since history has shown this latter issue to be a rather soft ground on which to built a useful discussion, we will approach the original question with respect to the potential of the neural networks approach to meet the engineering goals which are targeted by the AI research. So we have to look at what neural networks have accomplished with respect to such targets, what is their future promise/potential, and what is the implied impact.

Some of the major targets of AI research have been: pattern recognition/classification, automatic knowledge acquisition or learning, and associative recall. Of the many AI targets these are the most important, probably because of the relative difficulty of approaching them with the more traditional AI methods and models. There is not much need to expand on reviews of how successful the neural networks approach has been in these areas, the evidence is ample and well establishes the neural networks as a significant tool for approaching some of the major targets which have been set forth in AI research.

Neural networks have provided a way to approach AI targets on the basis of analytical models and methods. In neural networks we usually start with a computational structure which can be mathematically formulated and studied. With all due respect to many other AI methods we have to acknowledge that mathematical formulation is not always possible and the skill of the programmer is fundamental to their usefulness. In contrast, neural networks have made it possible to employ in the AI methodology some mathematical tools that were not employed before. They also made possible the transfer of methodology from other -seemingly unrelated- areas (such as physics, thermodynamics, systems control, signal processing etc.) in the AI arena. That this kind of transfer is made possible through the neural networks approach is not surprising; AI has been almost exclusively based on symbolic processing and first order logic. Neural networks on the other hand are akin to numerical estimators and dynamic systems. So it was natural that once a usefulness was established for neural networks, all the knowledge which had been accumulated about numerical methods and dynamical systems was automatically put at the same service. This bridge between AI and nonlinear dynamic systems has probably been the most important impact and the distinct contribution of the neural networks approach. With this bridge the AI research is equipped with an alternative to the traditional paradigm of finite state machines.

But this bridge is also the key to the potential and the future promise of neural networks with respect to AI. Traditionally whenever some goal of AI was approached, AI was then "redefined" and the new goals were moved further (e.g. game playing to problem solving to expert knowledge acquisition etc.). Now we are slowly but more and more targeting simulation of the cognitive functions of the brain. Neural networks have provided a basis for such simulations although to a limited extend. They have certainly demonstrated their inherent capability to perform generalization and abstraction which are distinct properties of higher-level cognitive processes. Neural networks are trained on a finite set of examples and then they are able to extrapolate on them in order to provide outputs to new inputs never encountered before. They also remain quite robust in noisy environments. They do not depend on the skill of some programmer. And after they are trained one does not really know what exactly they came to "know" or what they have actually learned. This is probably the major limitation right now but what is important is that it does not appear to be a dead end; an entire background in stochastic approximation, statistical methods etc. is there to support the study of their behavior.

Current neural network models are based on classical dynamic systems and therefore have in principle a deterministic behavior. The behavior of such systems is determined by the initial conditions. Variations in the behavior is due to sensitivity to initial conditions or to the problem of numerical precision but there is nothing inherently nondeterministic in their nature. Looking at neural networks from the computational capability (as in theory of computation) point of view, this inherently deterministic nature leaves neural networks with only the advantage of large memory due to their nonlinearity. But of course *good memories do not make thinkers.* However, the neural networks approach provides the base to employ non-Lipschitzian dynamics, that is systems with inherent nondeterminism, in AI. So even in the near future the neural networks approach will continue to break new grounds for AI.

Gerald Pechanek

I believe, artificial neural networks represent a means, i.e. a tool, for understanding aspects of artificial intelligence. A "good" method for understanding and explaining a concept is the use of analogies and alternative explanations. Artificial neural networks can be considered a massive parallel computational paradigm that offers new ways of viewing problems in artificial intelligence. For example, problem representation on the processing elements, need for high connectivity, requirements for precision of calculations, recursive nature of computation, learning, etc., are all problems addressed by artificial neural networks that have analogies in AI.

From a pragmatic "engineering" point of view, the development of a high performance neurocomputer requires the addressing of many of the difficult problems encountered in the design of massive parallel computer systems, such as computation and interconnection requirements. The solutions found into other areas of computer science, such as artificial intelligence applications and architectures.

Benjamin Wah

Artificial neural networks can be considered in a number of ways. They can be treated as a knowledge representation scheme, as a model for learning, as a mathematical model of computation, and as a model of the brain. The basic notion is distributed knowledge representation in which knowledge, at some levels of abstraction, is distributed in multiple storage elements, and each storage element is shared among multiple knowledge entities. In this sense, artificial neural networks can be treated as a way of representing knowledge in artificial intelligence. In a different perspective, they can be considered as a way of learning knowledge without first defining a specific representation scheme. This is most useful in problems in which the set of parameters is not well defined or the objective of the problem cannot be expressed in terms of measurable parameters of the problem. In this way, artificial neural networks are different from traditional machine learning methods that require some certain domain knowledge of the problem from the users before learning can be applied.

Learning speeds and performance of artificial neural networks can be greatly improved if they can be built using VLSI technologies or be implemented by fast parallel computers. Two major issues need to be considered in such implementations. First, hardware implementations improve the turnaround time of a given solution. They may not improve the "quality" of knowledge learned. Second, many applications are so complex to start with that any amount of parallelism available today may not solve the problem. For these applications, it is very important that the designers analyze their applications and decompose them into solvable modules.