Another View of Neural Networks

Cooper, Leon (1972 Nobel Laureate)

When interest in neural networks revived some fifteen years ago, few people believed that such systems would ever be of any use. Computers worked too well; it was felt that they could be programmed to perform any desired task.

Clearly, fashion has changed. Now, limitations of current computers in solving many problems involving difficult to define rules or complex pattern recognition are widely recognized; if anything, expectations for neural networks may be too high. The problem is no longer to convince anyone that neural networks might be useful, but rather to actually incorporate such networks into systems that solve real-world problems economically.

Neural networks are inspired by biological systems where large numbers of neurons, that individually function rather slowly and imperfectly, collectively perform tasks that even the largest computers have not been able to match. They are made of many relatively simple processors connected to one another by variable memory elements whose weights are adjusted by experience. They differ from the now standard Von Neumann computer in that they characteristically process information in a manner that is highly parallel rather than serial, and that they learn (memory element weights and thresholds are adjusted by experience) so that to a certain extent that can be said to program themselves. They differ from the usual artificial intelligence systems in that (since neural networks learn) the solution of real-world problems requires much less of the expensive and elaborate programming and knowledge engineering required for such artificial intelligence products as rule-based expert systems.

In their current state, neural networks are probably best at problems related to pattern recognition. Some existing neural network systems can efficiently and rapidly learn to separate enormously complex decision spaces. The problem of coordinating many neural networks, each a specialist in dividing some portion of the decision space, has also been solved. It is in these areas, therefore, that the first commercial uses will appear. Products that recognize characters, assembly line parts or signatures, that make complex decisions mimicking or improving on human experts (such as underwriters) that can diagnose engine or assembly line problems are in the prototype stage and/or are already fielded. One expects, further, that the pattern recognition ability coupled with, and feeding back and forth to rule-based systems (as has already been done in some simple applications) will finally result in machines that share our ability to learn and duplicate our processes of reasoning-machines that might be said to think.

The question is not whether but when.

Predicting the future, as we all know, is risky. Predicting the evolution of new technology is downright hazardous. Who in the 1930's would have said that among the consequences of the uncertainty principle would be transistors, silicon chips, and all of the vast array of solid state devices on which all modern computers depend? Or that superconductors would lead to extraordinarily sensitive detectors of magnetic fields now carried on many naval ships? Or in the late 19th century, that among the consequences of the research of Maxwell, Lorentz and Einstein, would be all that we call modern communication: radio, radar, etc.?

Accepting that risk, I would predict that neural networks will become standard components of what we today call computers. This will likely occur in a somewhat evolutionary manner: they will encroach gradually-board by board, intelligent components, that can be trained by humans in a language humans understand, into dumb machines-somewhat like neo-cortex came to dominate the reptilean brain. And, just as the 20th century is the century of automobiles, airplanes, telephones and computers, the 21st will be the century of intelligent machines. We will not only learn to live with these machines but, indeed, will wonder, one day, how we ever lived without them.