

THE LAST WORD

AI: great expectations

Artificial intelligence (AI) has repeatedly inspired great expectations in people who see the possibilities of applying its techniques. Sometimes it delivers. Sometimes expectations are dashed.


The idea that machines can be rendered intelligent has always been seductive, and demonstrations of limited scope tend to raise greater expectations than hindsight analysis shows were warranted. In his 1949 book *GIANT BRAINS or Machines That Think*, Edmund Berkeley ponders the amazing ability of machines such as ENIAC carrying out 500 multiplications of two 10-digit numbers per second, and envisions machines that would act as automatic stenographers, translators, and psychiatrists.

This pattern is still evident. A few years ago there were high hopes that robots would revolutionize factories. In a way they have, but not in the grand manner predicted in the business plans of start-up companies six or eight years ago. Programming problems, combined with a lack of flexibility, made it impossible to overcome the systems-level problems of integrating assembly robots into the world of manufacturing. Great expectations raised by demonstrations of such robots glossed over other critical aspects of a complete operational enterprise—not the least of which is capital. Robots, indeed, have proven to be useful, but not as useful as was first predicted.

We recently entered the bust side of another set of boom expectations: expert systems. Though mildly successful in industrial applications, once again the expectations of the expert systems industry have not been borne out over time. The too simple representations of the problem domains of expert systems make them extremely brittle when the world diverges from the narrow range of applicability of their knowledge rules. Broadening those representations has been frustratingly difficult. Expert systems are here to stay—though not yet ready to solve all of the world's problems.

With every bust there is a new boom, and in the past year just what the new fashion will be has become clear—neural networks. These networks incorporate an appealing idea in that instead of having to work out all of the details of a task, we'll simply let some randomly organized network of neuron models "learn" through trial and error how to do the right thing. Although neural networks have rarely accomplished anything beyond a computer simulation, business plans are being cranked out for new start-up companies to apply the technology.

But the current neural networks phenomenon is more than just another set of high expectations. This is the second time around for neural networks. It happened in the early '60s. In 1962, a distinguished Stanford professor predicted that computer programming would be obsolete by 1963 because, by then, users would simply converse in English with the front-end neural networks. Since then, there have been a few technical improvements, and computers are much faster, broadening the scope of the applicability and likely successes of neural networks. But, again, they really can't be expected to solve the world's problems. The old-timers, "immunized" the first time around, seem less enamored than the new converts.



I recently worked with a group from industry, offering a detailed explanation of a technical AI method. After some time, the lead technical member of the group—who had no previous exposure to AI—exclaimed, "But that's not intelligence! All you're doing is writing a computer program to solve the problem." Well folks, I'm sorry—but that's all there is. There is no magic in AI. All we do is tackle areas and tasks that people previously were unable to write computer programs to handle. Because we have developed sets of tools and methodologies throughout the years to accomplish this, AI approaches have had a number of good successes. But there is no universal set of magic ideas.

Every so often a new AI development comes along and great excitement ensues as people stumble over themselves, convinced that the key to intelligence has been unlocked. Maybe it *will* happen someday, but I rather doubt it. I don't think there is a single key to intelligence but rather that, unfortunately for both the philosophers and dreamers, intelligence is a vast, complex collection of simpler processes. To develop truly intelligent computers and robots, we're going to have to unlock those processes one by one—causing flurries of great expectations, followed by more modest real successes. This may sound boring and unimaginative, but I find it exciting. Intelligence really is a complex interaction of many things. As we unlock its secrets in the next few years and decades, we will see a constant flow of ideas that have real and immediate practical applications. Finally, when we truly understand AI, it won't seem like just a computer program but will appear as a wondrous testament to the creative genius of evolution. ■

Rodney A. Brooks

Dr. Rodney A. Brooks
Associate Professor
Electrical Engineering and Computer Science Dept.
Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, MA