How Long Will Computers Stay "Dumb"?

"It appears," writes Frederic Withington (1987) in a recent issue of the *MIS Quarterly,* "that information systems will always be what a child calls 'dumb'." I'm not so sure.

I share Withington's concern over much of the hyperbole that surrounds today's talk about artificial intelligence (AI). Many of those who try to sell us such supposedly intelligent merchandise as the expert systems or ordinary language interfaces seem (to me and Withington) to promise a lot more intelligence than they can actually deliver today. I am reminded of the story about the boy who cried "wolf." He was, if you will recall, wrong the first few times but right in the long run. The same scenario may happen with the AI people's cry of "smart" systems. Notice that I said "may" and not "will."

We often seem to be too optimistic when we predict what a new technology will do in the short run, but not optimistic enough in the long run. When the computer was first invented, for example, it was widely touted as a giant brain and there were predictions that it would be playing world championship chess by 1960. That short term prediction was, as we now know, a bit premature. Computers do not play chess at a world championship level yet. Some of the long term (and far fetched for the time) predictions turned out to be surprisingly conservative. Even though computers (based on vacuum tubes) were large, unreliable and expensive, a few optimists predicted as many as a thousand would exist by the year 2000. That turned out to be a "bit" on the low side.

Withington is, I agree, right in the short run. There are no really smart systems around today. In my personal opinion, the concepts used in systems being touted as smart today — the concepts that underlie the expert systems, the ordinary language interfaces and others — are too primitive to ever lead to smart systems. He is right in warning those of us who have to use them to beware of the accompanying promises. It is wise to assume that a system is not smart until you have tried it out yourself. Don't be misled by promises or canned demonstrations.

It might also be wise to keep our minds and eyes open. Something new that makes smart systems possible just might be invented. After all, analysts/technicians of the 1950s would probably have said that personal computers could not be built because the vacuum tube was the only thing available to build them. They would have been wrong, of course, because their thinking was constrained by the vacuum tube. But it took a Nobel-prize-winning invention — the transistor — to make them wrong.

I do not know what the basic breakthrough, or "transistor," of the intelligent machine will turn out to be, or even if such a thing will happen. Let me suggest some ideas. (The references to the literature are intended to be representative of each of the schools of thought I discuss.)

- **Parallelism:** We know that the components of the brain are much slower than the components of the computer and yet the human brain can do many things much faster than the computer can. Perhaps that is because the brain seems to work in parallel rather than serially. Traditional computers, based on what has come to be called the *von Neumann architecture,* can process information only one piece at a time. Parallel computers can process many pieces at once — for example all the pixels of a picture in one cycle rather than only one pixel at a time. Perhaps the increase in speed and the influence on the way we look at problems that parallel machines bring will be the "breakthrough" that will bring us smart machines by the year 2000 (Hillis, 1986). Perhaps not.

- **Symbiosis:** Perhaps the road to intelligent computers lies in the idea of symbiosis. Computers are smart in some ways and people in others. Computers can do routine work faster and more accurately than people. Humans can plan better and pay better attention to the big picture. Perhaps computers will only be able to increase human intelligence by adding their own peculiar abilities to ours. Perhaps tomorrow's smart information systems will have both people and computers in the loop.

  Computer systems for making medical diagnoses are a good example of such symbiosis. Such systems are seldom if ever asked to do full medical diagnoses. Rather, they suggest diagnoses. For example, instead of saying to the physician "Do so and so," they say "I think this patient may have A or B. My reasons for thinking she has A are so and so and my reasons for thinking she has B are thus and so." The computer has a better memory for many facts and more patience as it looks through those facts than the human physician may have. But it lacks the human's judgement. Together, they may make a smarter physician than either alone. A good example of such systems are those of Miller (1984).

Perhaps computers can only become as smart as they are efficient, but for some purposes that may be enough. For now, the only smart information systems I trust are those comprised of both humans and computers.

- **Learning:** Perhaps we cannot program computers to be intelligent. Perhaps we will have to train them, as we now have to train people. For instance, a new employee is trained to think as an employer would like him or her to think. Likewise, an employer might have to train a new computer. It would not be delivered already smart, just as people are not born smart. Perhaps we will have smart computers only when we learn how to teach them. We are beginning to learn how to let computers learn and that might be the road to the smart machine (Mitchell, 1986). Or it might not.

- **Brain Modeling:** Although the idea of developing an intelligent system by modelling the brain in a computer has been around for some time without achieving notable successes, a revival of this concept is occurring as our understanding of the brain grows. Perhaps the way to develop a smart machine is to imitate nature (Pellionsz, et al., 1977), although that did not work when we tried to make machines that fly. But one never knows.

- **Connectionism:** In the days of behaviorism in psychology and of the perceptron in computer modelling of the mind, information in the brain was thought to be represented by the connections between memory locations rather than as in today's computers, by the data stored in those locations. Today, a large active group of researchers at the University of California at San Diego, the University of Rochester and many other academic centers are reviving that idea (Rumelhart, et al., 1986). Perhaps. . .

- **Functional Programming:** LISP is widely used as the programming language for AI because it allows one to deal with both programs (which tell the computer what to do) and data (which tell the computer what to do it to) in a uniform way. The ability to treat programs as data allows programs to change themselves and to "think about" themselves. Functional programming allows us to fudge this distinction even further and people at MIT are excited enough about this idea that all students majoring in computer science take a course on functional programming as their introduction to computing (Abelson and Sussman, 1985). Perhaps. . .

- **Noncomputations:** All the ideas I have suggested so far stay within the realm of computations. But it is becoming more and more widely recognized today that the machines we call "computers" can actually carry out processes that are not, technically speaking, computations at all. And such processes can do things that computations cannot. Perhaps intelligence lies within their domain (Kugel, 1986). Perhaps not.

- **None of the Above:** Perhaps the breakthrough will come in some area that I have not included in my list. That would not be surprising since it is the nature of breakthroughs to be surprising.

Withington seems to feel that computers cannot be "smart" because they can only do what they are told to do. They do not, he argues, understand why. They lack consciousness, volition or purpose. Withington may be right; but he may be wrong for two reasons. He claims that humans actually have these special features; however, as we have learned, we can be mistaken about how our own minds work. We may think we have purposes and self-consciousness, but perhaps what seems to us to be self-consciousness or purpose, is really something else. If it was, it would not be the first time that introspection had been wrong. We trust introspection because our mind tells us we are right. But our accounts of how the mind works have often been wrong as they might be here.

Withington may also be wrong in his claim that computers cannot have these features. Many of the approaches to AI listed above are based on theories of how computers might be given purpose, consciousness, and the like. And those theories might give us ways to develop these "human" features into machines.

I am not as sure as Withington that information systems will stay "dumb" through the first part of the twentieth century. Let us, by all means, be careful about all the rosy claims we hear of intelligent machines today, but let us also remember the boy who cried "wolf." He did not stay wrong and those who cry "smart machine" may not stay wrong either.

Peter Kugel
Boston College

References

Abelson, H. and Sussman, G.J. *Structure and Interpretation of Computer Programs,* MIT Press, Cambridge, MA, 1985.

Hillis, W.D. *The Connection Machine,* MIT Press, Cambridge, MA, 1986.

Kugel, P. "Thinking May Be More Than Computing," *Cognition,* 22, pp. 137-198, 1986.

Miller, P.L. *Critiquing Approache to Expert Computer Advice - ATTENDING,* Pitman, Boston, 1984.

Mitchell, T.M. *Machine Learning: A Guide to Current Research,* Kluwer Academic Publishers, Boston, MA, 1986.

Pellionsz, A., Llinas, R. and Perkel, D.H. "A Computer Model of the Cerebellar Cortex of the Frog," *Neuroscience,* 2, 1977, pp. 19-36.

Rumelhart, D.E., McLelland, J.L. and the PDP Research Group. *Parallel Distributed Processing,* MIT Press, Cambridge, MA, 1986.

Withington, F. "The Mature Intelligent Computer," *MIS Quarterly,* 11:1, March 1987, pp. 1-3.