J. David Bolter

Artificial Intelligence

ARTIFICIAL INTELLIGENCE, writes Marvin Minsky, is "the [redacted] erational planning so popular with business and the military will be performed entirely by computer.[2]

For years we have been reading in newspapers and popular magazines about the wonders of computers, and for decades, science fiction movies and books have featured all-knowing, though often malevolent, computers and robots. Many laymen believe that artificial intelligence is already with us, that some computers can understand the spoken word, read the newspaper, pilot a car using a television camera, or design their own even more talented offspring. In fact, no current program can do more than parody these achievements. Artificial intelligence is much harder to come by than the real thing: excellent programmers have been working for years on such problems and have made only a little progress. There is an enormous gap between what computers were built to do (mathematics and symbolic logic) and the wide range of skills that humans possess. Programmers must still work close to their machine's natural talents.

Computers are good at such well-defined and logical games as checkers, backgammon, and chess. They can play the first two about as well as any human, and can defeat any amateur at chess. For ar-

There are indeed practical applications for the techniques of arence of the programmer. One program may fly the space shuttle from orbit to landing with no help from the pilot; another may type out a letter in perfect paragraphs with underlining and boldface for emphasis. A program can even surprise its creator, often by committing some disastrous error. This ability to dispense with human control while performing complex tasks makes the computer metaphor almost irresistible to millions of professional and casual programmers. Is there not something human about this machine? Or is turn of the century, the prediction will probably fail. The terms of the test, however, are still quoted with approval by those working on artificial intelligence, and they remain faithful to their goal of proving that a computing machine can successfully imitate its creator.

In this respect, the artificial intelligence project is part of a long tradition.[6] For centuries, engineers and craftsmen have been making machines that imitate human beings, although their efforts have never been taken as seriously as the current project. The immediate forerunners of intelligent computer programs were electromechanical devices. In the fifties, for example, Grey Walter, a physiologist, designed an electromechanical "turtle" that could negotiate its way around obstacles on the floor. The turtle's technological ancestors include an electromechanical chessplayer built by the Spanish technologist Torres Quevedo. Before that, there were generations of fully mechanical automata. The most famous in the eighteenth century were those of Jacquet-Droz, including a boy who could write a message with pen and ink and a girl whose moving fingers could play a tune on a harpsichord, and a duck by Vaucanson that could quack and even eat and eliminate its food. The gardens of Europe in the seventeenth and eighteenth centuries were also adorned with figures mechanical devices could only point the way to the fully electronic technology that can express human nature.

There is some justice to this claim, because the computer *is* a more [redacted]

[redacted]

[redacted] metaphor. The idea of comparing the bodies of humans and ani- [redacted]

Load an artificial intelligence program into a computer and set it running. Of what is the computer now a model? The term "electronic brain" was once a common expression of the computer metaphor, but is the computer a model of the human brain? This was proposed in the early days of computing. Turing and another pioneer, John von Neumann, had thought of comparing the brain's "hardware" with the computer's. As early as the 1940s, Warren McCullough and Walter Pitts had described the mathematical properties of a net of neurons, drawing an analogy between electronic components and the human nerve cell. At the time, neurons still seemed to be fairly simple processors of electric information, and a digital computer might hope to reproduce the network of neurons in the brain. But the hope was never fulfilled. Although some elegant mathematics—was discovered along the way—the elementary brain machines themselves were never more than toys. Scientists also began to realize that the brain, with its forests of dendritic connections and complex chemistry, was not a simple network of discrete logical experience, allowing scientists to predict and control events—that, as Galileo said, the book of nature should be written in the language of mathematics. The mystery of the mathematical prediction and control of nature remains a subject for philosophers of science. My

gence are not like the mathematical models of physics. Computer models are not analytic in the same sense; they do not rely on deep mathematics, and they have no vitality and carry no intellectual conviction when separated from their machine.

Consider how the physicist and the artificial intelligence programmer each use the computer in their work. For the physicist, the equations and their interpretations form an abstract world that he can ex-

The artificial intelligence specialist would say that this program "understands" Chinese, for he accepts the definition that to know Chinese is to be able to transform input sentences into meaningful output. But Searle replies that the computer so programmed does

[redacted]

ment by way of proof. Reverse the customary process of automation: replace the computer with a human being, one who speaks English and knows not a word of Chinese. Write out all the rules of the program in English, and let this human information processor apply these rules to sentences given him in Chinese. If necessary, have him memorize the whole procedure. Now this person can read stories and answer questions in Chinese, totally mechanically and by the way quite slowly, but he still cannot speak or write any Chinese

[redacted]

is all the computer has for anything it does. To confuse simulation with duplication is the same mistake, whether it is pain, love, cognition, fires or rainstorms.[9]

Here, the artificial intelligence project has indeed abandoned common sense. In its excitement over the computer metaphor, it must abandon common sense, and indeed all computer programmers are tempted to do the same. Every one is inclined to confuse simulation with duplication because of the protean nature of the computer itself. This abstract machine can imitate the design, if not the physical effect, of any other machine and of many aspects of the natural world as well. The computer can simulate the operation of a jet engine, the traffic flow of a city, the reproductive rates of a colony of bacteria, or the invasion of Western Europe by Soviet tanks. Any

phenomenon that can be divided into a series of discrete events can be simulated with some success by the computer. Every significant program is a simulation, an attempt to recreate within the computer promoted by simulation, by trial and error, by thinking in terms of the process and its results rather than the deep causes. It may be that artificial intelligence will win this philosophical debate simply by sweeping away its opponents, refusing to engage them on their own ground. It will then offer its own answer to the great philosophical question: What is knowledge? The answer—that knowledge is the manipulation of symbols according to formal rules—will be wholly

unsatisfactory to analytic philosophers and indeed to many philosophers of the older schools. The whole question will have been recast

strings? The bits must be given a structure, a particular order and context that make them meaningful. Indeed, determining the appropriate data structure, deciding how the computer will interpret its ones and zeroes, is the programmer's main task. This is as true of the mathematician and the city planner as of the artificial intelligence specialist: they must all find suitable ways of arranging their data. However, the artificial intelligence programmer has the unusual task of finding data structures that will reflect the interplay of

Consider for a moment the differences between the psychology of artificial intelligence and psychoanalysis, perhaps the most influential psychology of the last hundred years. Psychoanalysis is analysis in the old style: its seeks to probe beneath the surface, to find deep causes behind human actions. Its metaphor of the mind emphasizes

the id, buried below the more rational layers of ego and superego. The goal of psychoanalysis is to expose the repressed memories of childhood, and the psychoanalyst explains the human mind not as a

The more the Freudian psychologist probes the mind, the more he realizes the endless character of his task. But the specialist in artificial intelligence does have the experience Minsky describes: he does "reach bottom," almost immediately, as he maps complex human [illegible]

The artificial intelligence specialists have, I think, gone too far. The computer is a mirror of human nature, just as any invention reflects to some extent the intellect and character of its inventor. But it is not a perfect mirror; it affects and perhaps distorts our gaze, magnifying certain human capacities (those most easily characterized as "infor- [illegible]