Bridgewater Review

Volume 3 | Issue 1 Article 7

12-1984

Artificial Intelligence: Myths and Realities

Hugo D'Alarcao
Bridgewater State College

ARTIFICIAL INTELLIGENCE

MYTHS and REALITIES

by Hugo D'Alarcao

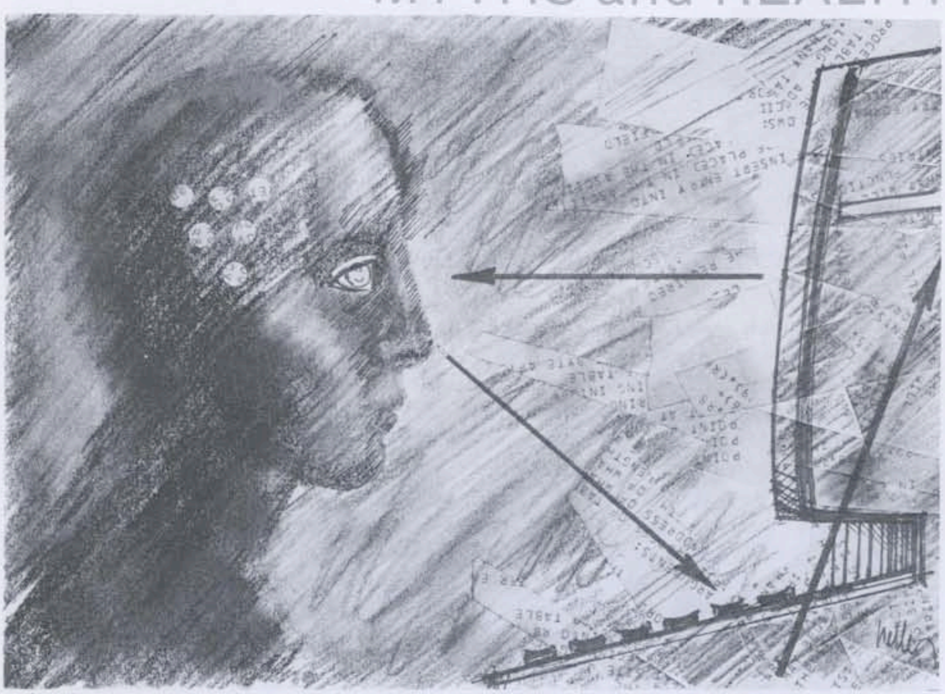

Illustration by John Heller

Artificial intelligence -- the name conjures images of mechanical monsters, the Golem, Dr. Frankenstein's creation and the rebellious computer Hal. We have always been fascinated by the possibility of creating a machine in our image, but this fascination is often accompanied by apprehension. We fear losing control of our creation and suspect that it might turn against us. It is this duality, this conflict between the desire to create and the fear of the consequences of the creation, that has been so successfully exploited by writers. It is also, in part, this fascination that has recently brought the field of Artificial Intelligence (AI) into public view.

Some recent developments with far-reaching practical applications have shown the enormous potential of AI. If we look beyond the spectacular and leave behind our apprehensions, illusions, and fantasies, we may sense the real and lasting significance of AI. To help dispel some of these myths, I will describe briefly the goals and the major subdivisions of AI.

The first myth is that AI is concerned with the question of whether machines can think -- it is not. The question may be an excellent topic of conversation at parties but is of no interest to practitioners of AI. Probably the best way to describe Artificial Intelligence is to specify its goals. There are two major goals of AI: to make computers more useful, and to understand what intelligence is and what makes it possible. These two goals define, to a great extent, the different branches of AI. The primary goal of *expert systems, natural language processing, vision,* and *robotics* is to make computers more useful. By contrast, *cognitive modeling* and *machine learning* are primarily concerned with understanding the possibilities of intelligence.

Expert Systems

The most spectacular successes of AI have occurred in the field of expert systems. An expert system is a computer program that simulates the expertise of a specialist in some field. One of the best known of the expert systems is INTERNIST/CADUCEUS at the University of Pittsburgh. It covers more than eighty percent of all internal medicine and diagnoses at the level of a medical expert. Another program, MYCIN, at Stanford University diagnoses and recommends treatment for infectious blood diseases. When a panel of human experts evaluated and compared the performance of medical experts, interns, and MYCIN, the computer program's performance was judged as good or superior to all the others. PUFF is a computer program to diagnose pulmonary diseases. It is now routinely used as a consultant at the Pacific Medical Center in San Francisco. These and other medical diagnostic programs have shown that expert systems are feasible and perform at or above the level of a human expert.

Currently these are very large programs requiring the huge capacities of large computers called main frames. It is, however, realistic to think that in the near future similar programs will be available for microcomputers. When this becomes reality all those regions of the world where medical care is either non-existent or minimal may have easy access to the best diagnostic facilities.

Another famous expert system, and one of the earliest ones, is DENDRAL. This system, at Stanford University, has been evolving for over sixteen years. It analyzes mass spectrographic and other chemical experimental data to infer the plausible structure of an unknown compound. DENDRAL by now surpasses all humans at its task. This sytem is also of interest because it typifies the kind of cooperative efforts required for the creation of expert systems. DENDRAL was started when Edward Feigenbaum, a computer scientist, met Joshua Lederberg, a Nobel laureate in genetics. Together they formulated the idea of a computer program to infer molecular structure from chemical data. Together with Carl Djerassi, a physical chemist, they created DENDRAL. It has been growing in sophistication and scope ever since, and is now used at university and industrial chemical labs throughout the world.

In biology there is a program called MOLGEN (MOLecular GENetics) that acts as a consultant in genetic engineering and analyzing DNA sequence data. Computer manufacturing also benefits from expert systems. Digital Equipment Corporation uses an expert system to design how the

different modules of VAX computers are to be put together to conform to specific customer needs. This system is reported to plan correctly in more than ninety-nine percent of the cases (far better than the human counterparts!)

One spectacular success was achieved by PROSPECTOR, an expert system used in exploration for minerals. In 1982 this system made a find of a molybdenum deposit estimated to be worth up to $100 million. This find was missed by the human experts, even though they had considered the site.

There are many other examples of expert systems, now the most active field of AI. One of the most interesting conclusions to come out of the study of expert systems is that knowledge-intensive fields such as medicine lend themselves most readily to simulation by computer programs. The hardest thing to simulate is everyday reasoning! Another important by-product of the study of expert systems is the recognition that the hardest part of translating a human expert's knowledge into a computer program is obtaining the needed knowledge from the human expert. This realization, and the accumulated experience of twenty years of study of expert systems, has given rise to a new discipline: knowledge engineering. Expert systems, because of their spectacular successes and obvious applications, have stimulated activity in other branches of AI. For example, given the time consuming and difficult task of obtaining the necessary knowledge from human experts, it would be very desirable to have the system learn the needed knowledge by itself from books or human teachers. This desire has stimulated research in the field of machine learning. Similarly, since it is important to facilitate the communication between the expert system and the user, there is a need for natural language processing.

Robotics

This is another field whose primary objective is to make computers more useful. Currently the main use of robots is as manipulators to perform industrial tasks such as assembly, welding, painting, and other repetitive or hazardous tasks. No, robots do not look lik R2D2! Most of them are mechanical arms such as the arm in the space shuttle. At this time most industrial robots have a limited capacity to perceive or to respond to their environment. A robot must "know" where it is in relation to the objects it is to manipulate. However, in university and industrial labs, prototypes already exist that are somewhat autonomous in their interaction with their environment. Although several companies are already marketing home robots, the time for their widespread use has not yet arrived. Most robots are being used for industrial applications, for tasks such as material handling including loading and unloading, moving parts, assembly tasks, spot and arc welding, and painting. Several factories have already been fully automated by robots and many others are now partially mechanized. The rate of expansion of robot use in industry may be measured by the amount of robot sales. In 1976 robot sales were $15 million. By 1979 it was $45 million. In 1983 sales reached $250 million, and it is predicted that by 1990 it will reach $2 billion. This rapid automation may produce problems not too dissimilar from those of the industrial revolution. Much more attention needs to be given to the social consequences of the widespread use of robots. Such automation, if it is carried out at a rapid rate and without adequate planning, may bring about widespread unemployment and the need for large retraining programs. Even in the best of all possible situations such dislocation of jobs will increase the leisure time of most people and this, in turn, will require careful planning so as to avoid alienation, boredom, and other individual and social ills.

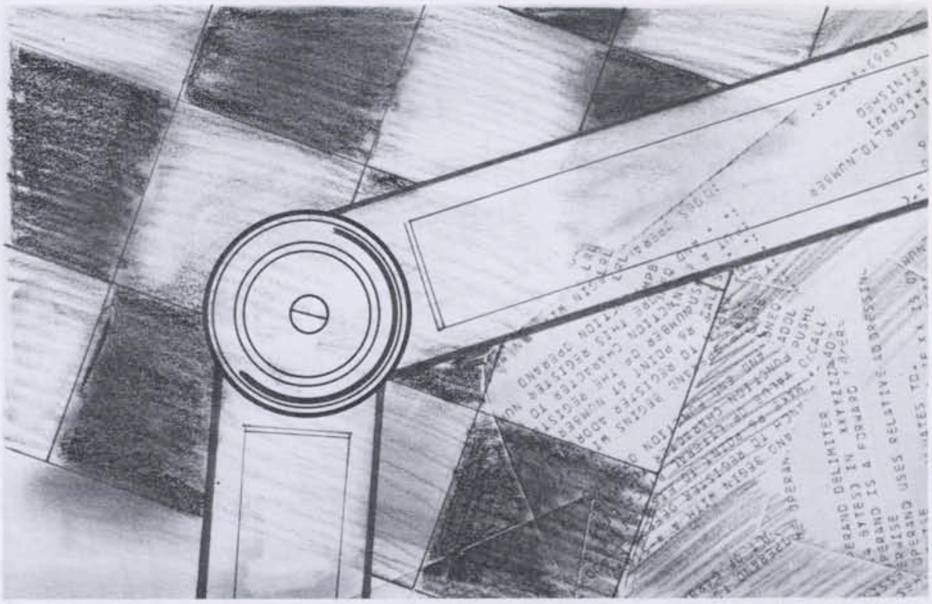

Illustration by John Heller

Robotics research encompasses many different fields. Aside from all the engineering aspects of robotics, research is being carried out by mathematicians, computer scientists, and neurophysiologists. For example, the motion of robot manipulators requires sophisticated mathematical techniques. Let us look at a typical problem, motion planning. Suppose that the geometry of a robot's environment is totally known. That is, we know not only the geometry of the manipulator itself, but the shape and location of all the objects in its working space. If we are given an initial position for the manipulator and a desired final position, how do we plan a continuous motion that will take the manipulator from its starting to its final position avoiding any obstacles along the way? Even if we succeed in solving this problem abstractly, how do we translate it into a sequence of executable motions? How do we control the actual motion of the arm? How can we adjust for errors due to mechanical limitations? These are only some of the very difficult questions that workers in manipulator motion are

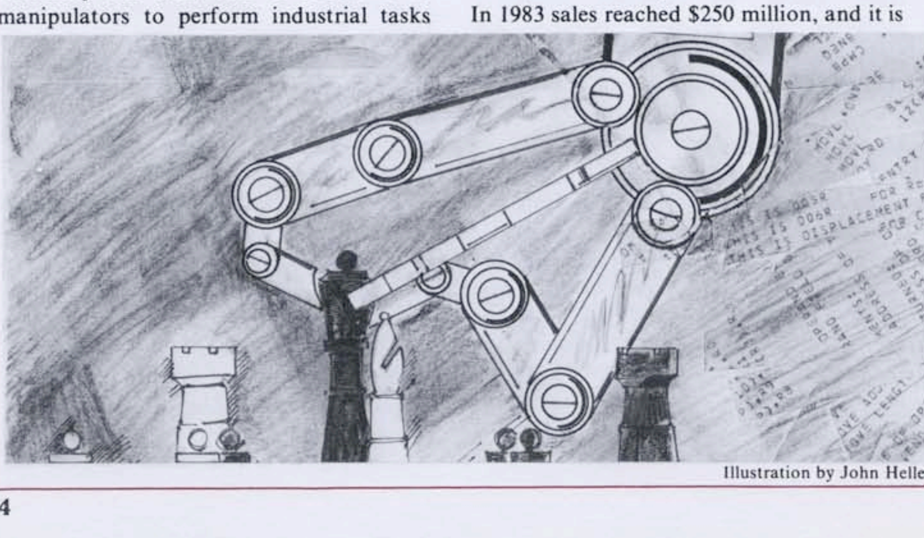

Illustration by John Heller

considering. Related questions concern the efficiency of the procedures for motion planning and control. Not only must we have procedures to control and plan the motion of the robot arm but they must be efficient. They must be executable in a reasonable amount of time. Imagine the problems we would face if every time we used our arm to reach out and grasp an object our nervous system would require a few minutes to plan and control our arm's motion!

Another very active area of robotics research is locomotion. Most current industrial robots are mechanical arms with a fixed, immobile base. One of the most important future applications of robotics is in the exploration of, and work in, environments that are inhospitable for humans, such as space, underwater, and hazardous industrial environments. The robots required for such tasks must be not only mobile but able to traverse very diverse types of terrains. The researchers concerned with these problems work closely with physiologists and zoologists to try to understand human and animal locomotion. For example, the latest issue (vol. 3, no. 3) of the *International Journal of Robotics Research* is totally dedicated to legged locomotion. A brief description of some of the articles in this issue may give the flavor of the work being done.

One paper studies the mechanics of human walking and running. The conclusions reached have led to the design of a "tuned track" where runners can compete with fewer risks of injury and faster times. Another article describes a 1600 lb., gasoline-powered, hydraulically actuated, six-legged walking machine that can carry a man. One paper quantifies the gaits of two and four-legged animals and concludes that the gaits of animals seem designed to minimize unwanted movements and energy costs. Another describes a one-legged hopping machine constructed at the robotics institute of Carnegie-Mellon University. Another article studies the tactics used by locusts for movement on rough terrain. Such basic issues of physics, biochemistry and other allied sciences must be understood if robots are to be successfully programmed to work in varied environments.

Another important area of research is the study of the manipulator "hand." The problems being considered include determining the optimal number of fingers for a given task and the desirable number of joints per finger for stable grasping of objects. Perhaps the most active area of research in robotics deals with perception. Computer vision, for example, is now an

Illustration by John Heller

important field in its own. Very deep understanding of the computational problems involved in both human and machine vision have been achieved by the late David Marr and his group at MIT. Their research has not only provided a theoretical foundation for the study of vision but illustrates how fruitful the collaborations between computer scientists, mathematicians, and neurophysiologists can be. There are presently several sophisticated vision systems being used by robots for tasks such as quality control of manufactured parts, analysis of printed circuits, and recognition of parts in assembly operations. Other perception studies involve tactile sensors being developed for manipulations and those which deal with how the senses of hearing and smell are used by humans and animals and what comparable mechanisms might be useful for robots.

The hardest thing to simulate is everyday reasoning!

One problem of great importance for industrial applications of robotics is the development of programming languages for robots, that is, languages specifically tailored for the programming of tasks such as motion planning and locomotion. Finally, one other problem has come to light because of the use of robots. It was discovered as soon as robots began to be used industrially that the assembly line, in particular, and the factory, in general, were designed for humans, and that considerable changes were required in order to make them efficient for robots. The automated factory is very different from the non-automated one; it requires a totally different design. This is yet another field of investigation.

Natural Language Processing

One of the earliest goals of AI was to have the communication between humans and computers in a natural language such as English. To date, this goal has not yet been reached but great progress has been made and important difficulties have been identified. Several systems currently exist that process natural language in specialized areas. Some of these are used in conjuntion with expert systems to facilitate the use and the growth of these systems. Another early goal was to have machines translate from one natural language to another. Progress in this field has been disappointing. The difficulties encountered were much greater than originally expected. A classic example of the inadequacies of machine translation is the following: after translating the sentence "The spirit is willing but the flesh is weak," into Russian, the computer then translated

the Russian back into English as "The vodka is good but the meat is rotten." Although natural language processing might be one of the most difficult of the problems tackled by AI, it probably will provide more insights into the nature of intelligence than any other. This brings us to the second goal of AI -- the understanding of the possibilities of intelligence.

Illustration by John Heller

Cognitive Modeling

Recently, linguists, psychologists, neuroscientists, and computer scientists have joined forces in an interdisciplinary effort. It was realized that by bringing to this collective effort the insight and techniques of the different specialties more than the sum of the parts could be achieved. The catalyst was the emergence of Artificial Intelligence as a serious science. The ultimate goal of this new breed of scientists is to explain, using computer programs as models, every aspect of cognition. The common assumption underlying this enterprise is that human beings and computers are examples of physical systems that hold and transform symbols. This new science is known as cognitive modeling. What makes cognitive modeling different from each of its constituent disciplines is the use of a computer program to model the theory. As anyone who has written even the simplest program knows, it is necessary to fully understand what one wants the program to do to be able to write it. The old truism that to determine whether you fully understand something you should try teaching it to someone else can now be considerably strengthened by making the test instead to try teaching it to a computer, that is, to write a program to carry it out. By designing a computer program to simulate a particular cognitive function it becomes more apparent what the limitations of that function are. That is, what it can and what it cannot carry out. As David Marr wrote, "The best way of finding out the difficulties of doing something is to try to do it."

It is in this respect that I see the great potential of AI. The type of phenomenon that sciences such as physics and chemistry study are, somewhat, stable. Even though our understanding of physics has changed considerably since the time of Newton the physical universe that Newton was interested in has not changed much, at least not on a local scale. A chemical experiment carried out today can be reproduced tomorrow. Because of this stability mathematics is the vehicle for the formulation of theories of the physical sciences. For example in physics, the most mathematical of the sciences, $E = mc^2$ is itself the embodiment of the theory. It is the statement of a physical reality that cannot be stated better any other way. Thus a science is mathematical if its truths can be embodied in mathematical facts. This has not been possible for the behavioral sciences. It is my belief that the reason that these sciences have not been able to use the language of mathematics, is that mathematics is not suitable to express their truths. However, I believe that AI is the vehicle that will do for the behavioral and social sciences what mathematics did for the physical sciences. It is computer programs of this type being developed by the cognitive scientists working in AI that will help model the truths of psychology. Indeed it is interesting to note that individuals such as Roger Schank at Yale and John Anderson at Carnegie-Mellon who are pioneering this type of work in cognitive modeling are psychologists as well as AI scientists.

Machine Learning

Another very active field of AI is the design of programs that learn. The efforts in this direction have not only led to interesting successes but have continually provided deep insights into the human learning processes. I will mention only a few such programs. AM is a program written by Douglas Lenat at Stanford University. This program started by assuming a few basic concepts of set theory as well as some way to measure what might be mathematically interesting. It then proceeded to discover most of the facts of elementary arithmetic. It began by discovering the concept of number and from there it came up with the arithmetical operations. It discovered prime numbers. It conjectured the fundamental theorem of arithmetic and rediscovered a famous mathematical hypothesis that has so far, defied verification or disproval.

Another learning program called BACON discovers empirical laws by detecting regularities in data supplied to it. It has rediscovered most of the early physical and chemical laws. Some of these include Ohm's law for electric circuits. Archimede's law of displacement, Gay-Lussac's combining volumes, Cannizzaro's determination of the relative atomic weights, and Proust's law of definite proportions.

ACT is a program developed by John Anderson to prove theorems in high school geometry. This program exemplifies what I mentioned earlier about the use of AI as the tool of cognitive science, it represents and tests Anderson's theory of learning as applied to the learning of geometry.

I have mentioned only some of the major areas of AI. Some of the techniques that AI scientists have developed have far-reaching consequences. For example, there is a new programming language PROLOG which may change our thinking about programming. PROLOG is an acronym for PROgramming in LOgic. This language , unlike all preceding ones, does not require the programmer to spell out in detail each of the instructions the computer is to carry out. Instead it requires that the assumptions and the desired conclusion be very specifically stated and leaves the way of logically deriving the conclusion from the assumptions to the computer.

This relatively new field of AI will surely bring about many new ways in which to think about thinking. It is natural that AI scientists should be excited by their achievements and revolutionary goals. However, it is important that we do not confuse excitement with immediate realization and that we prepare for the social impact of the successes of this most challenging and intellectually stimulating science.

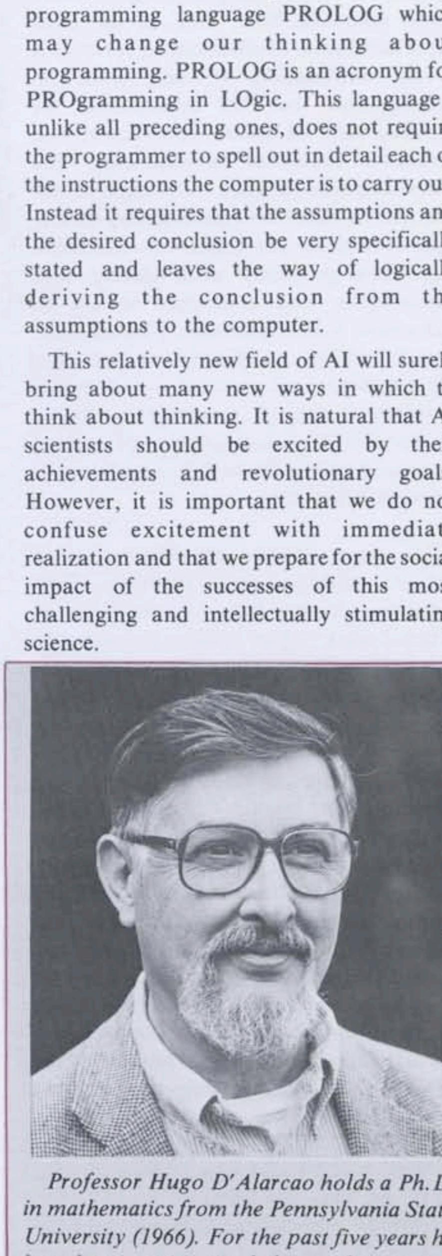

Photograph by Robert Ward

Professor Hugo D'Alarcao holds a Ph.D. in mathematics from the Pennsylvania State University (1966). For the past five years he has been active in helping design and implement the computer science program at Bridgewater State College. His current research interests are in robotics and computer vision and he is in the process of developing a robotics laboratory at B.S.C.