

GEORGE DRAPER  
JAN. 1985

Artificial Intelligence ,or  
Thinking about Thinking

We members of the Athenaeum Society are brought together primarily because of an interest in, and mutual respect for, intelligence, and the literary products of intelligence.

We think. It is such a simple, such an easy and pleasant thing to do. We read; we study; we speak; we act, and all with an ease that bespeaks simplicity rather than complexity underlying these performances.

Surely anything so easy to do as thinking should be readily outlined, reduced to some sort of system or pattern and, given the tremendous new powers of computers, be written in workable software form. That having been accomplished, the results could be crystallized into good workable hardware.

At least, so it would appear.

Let us examine briefly one of the most exciting projects man has ever proposed for himself, the construction of a dry brain, a device which can, like the wet brain of man himself, actually think, reason, solve complex problems, make judgments and, perhaps, in addition, even improve and eventually reproduce itself.

In this examination one may soon find that, in some ways, artificial intelligence is more closely allied to philosophy or psychology than to circuitry and electronic engineering.

This is a new field. The very name "artificial intelligence" was invented in 1956 by John McCarthy to describe a summer workshop at Dartmouth College.

The assumption that AI was even a remote possibility rested on a theoretical base set forth about a century earlier. To gain perspective let us briefly sketch some developments, beginning in the last century, which brought us up to John McCarthy and his workshop.

In the year 1854, Boole, in his published work "Laws of Thought," undertook to outline and construct "the mathematics of human intellect"--a formal logic distinguished by its ability to be represented in pure mathematical symbols. Boole stated and contended that the complete body of logic could be distilled, reduced to, and represented by nothing more than a series of yes-or-no responses, which in turn could be coded or designated in binary terms. Simply described, this was an algebra of thought, based solely on the numbers 0 and 1.

Innocent though this may appear, Boole's new theoretical approach to logic representation was the intellectual insight that initiated mental explorations which eventually culminated in the invention and construction of computers. Without Boole's insight there would be no computers, and no possibility for the creation of artificial intelligence.

His concept lay dormant for many years, a matter of slight theoretical and academic interest. Not until the 1930's did theory begin to gel and assume practical shape.

The computer industry owes much to Claude Shannon, who, as a graduate student at MIT, submitted as his master's thesis a paper with the title "A Symbolic Analysis of Relay and Switching Circuits," which, in a flash of brilliance, linked three things: binary math, symbolic logic, and the behavior of electronic

circuits. He offered in his thesis the startling, practical, and inspired deduction that the propositional calculus of symbolic logic could be used to describe the two-state, on-or-off behavior of an electromechanical relay switch.

From Claude Shannon's inventiveness came what is termed information theory -- the assumption that information, just like energy or matter, is an entity which is both quantifiable and capable of being manipulated at will.

Alan Turing, a Cambridge mathematician, visualized in the 1930's an imaginary computer--now termed a "Turing machine," essentially a machine with an endless tape, marked off in blocks of two types (like binary code). With four simple operations this machine had universal capabilities. This was the theoretical model of the computer to come.

It remained for Alonzo Church, a logician, to come forth with what we now call Church's thesis. His proposition was that a computer could be built that could do anything for which an algorithm, i.e., a precise and unambiguous set of instructions, could be written. Computer science rests firmly on this proposition, which constitutes in addition the theoretical ground for artificial intelligence.

The necessary theory was in place. Next came rapidly improving successions of hardware.

Not too long ago Intel founder Robert Noyce, writing in Scientific American, compared today's \$300.00 personal computer to the gigantic ENIAC of the 1940's : the home computer "is twenty times faster; has a larger memory, consumes the power of a light bulb rather than that of a locomotive, occupies  $\frac{1}{30,000}$  the volume,

and costs  $\frac{1}{10,000}$  as much."

These things are mentioned because, in the field of artificial intelligence we are confronted by the necessity of duplicating an enormously complicated natural computer with a hundred billion cells, with each of these cells making hundreds of dendritic contacts.

Consequently, even with the speed and capabilities of computers, doubling every two years for the last twenty years, (and this rate will no doubt continue), we have only begun to get computing devices having sufficient capacity to make raw beginnings in mimicing human intelligence.

In 1956 just about everything appeared possible. With four kilobytes of memory, AI programs of a sort could be written. The goal set was to duplicate the entire range of human problem solving abilities, and with Carnegie-Mellon University and Rand Corporation cooperating, in the year 1957 a program was developed called the General Problem Solver (GPS).

In the very simplest of cases, GPS worked. Confronted by more difficult problems GPS failed, and failed miserably and spectacularly. Subsequent similar programs also failed. It was determined that these programs failed because they were completely overwhelmed by what was termed the "combinatoric explosion." (too many possible combinations). To visualize this difficulty consider the GPS problem in chess. There are, unbelievably,  $10^{120}$  sequences of legal moves possible in the game. Should one wish to devise a totally general program, one that would examine every possible move, billions of years would be required to prepare the needed program.

Confronted by the enormity of the problem, AI programmers retreated from any pure general problem solving type of program and instead limited the scope of decision to narrow problems. "Heuristics," or rules of thumb, were employed in specific situations. Again relating this to the game of chess, the IF-THEN instruction would be given as follows: " If a pawn has a choice of moving forward or capturing your opponent's queen, then take the queen."

By the mid 1970's the whole idea of the general problem solver had begun to appear rather hopelessly naive.

A.I.'s new approach was to have available tremendous amounts of highly specific knowledge about a great variety of things. The programmers soon found that even "knowledge" itself is an evasive concept, and such "knowledge programming" was destined to be both frustrating and baffling. Still, there was success. The success was confined to operations within very clearly defined and narrow scopes or domains, for example medical diagnosis. The input in such a program would consist of careful and detailed information on symptoms, laboratory findings, history, etc., with "knowledge" being the programmed wisdom and experience of the finest diagnosticians in certain fields. Usually respectable diagnoses were produced. The emphasis here should not be on success, but on the necessity for operating in a highly restricted domain, or area of application.

There are difficulties in programming "knowledge". These can be divided into three categories.

First, to parallel human memory, there is the representation of knowledge.

Second, to parallel human problem-solving and planning, there is the control and use of knowledge.

Third, to parallel human learning, there is the acquisition of knowledge.

Consider the first item, representation of knowledge.

In programming it is not adequate simply to list names and numbers, putting together a string of data like the information in a telephone directory, for that is not knowledge. Knowledge must help the computer do smart things. In this connection Shakespeare's remark that brevity is the soul of wit is pertinent. The word "wit" meant "knowledge", thus Polonius was giving a formula for intelligence: the whole method of intelligence is to abbreviate and manipulate knowledge. A large knowledge base can become nothing more than an unstructured listing of statements, incomprehensible and useless to the programmer, for it neglects this one extremely important fact: our human knowledge is highly structured, and related by one thing suggesting another.

How do we structure and relate information? One approach is that of Scripts, pioneered by Roger Schank of Yale University. These are elaborate data structures that model human expectations. For example, the Athenaeum Meeting experience script would feature meeting fellow members, engaging in conversation, standing, engaging in conversation, seated, eating, listening to minutes, then listening to some tired, hopeless excuses for absences, and all the rest of the script: a framework of first Thursday expectations to aid in inferences.

One encounters a key problem with any rigid representation of knowledge. Humans are not mechanical and rigid. Faced by an

extremely complex problem, the human, initially baffled, can mentally step back, assume a different mental posture and, in many cases, solve the problem at once. For the human the solution is obvious. It follows logically that our human "knowledge" representation" is both high-structured and wonderfully reshapable: it is certainly not rigid.

Inference is buried in the AI programs as heuristics, deep in the code. It is efficient but rigid. Back in 1959 John McCarthy pointed out that problem-solving strategies would be far more flexible, and easier to understand if they were first dissected out and made explicit. It is too bad that this approach leads us straight back to the same quandry of the general problem solver and our old friend the combinatory explosion.

Humans employ nonstandard logic. Consider the difficulty of making a machine reasoning process automatic when the data base contains statements of the type "He believes this," or "She knows that"? Marvin Minsky emphasizes the problem of non-monotonic logic, giving particular stress to the ability we humans have to retract a conclusion after being given different information, an ability ignored by the heavy-handed AI focus on logical deductions. This AI problem is recognized as the Minsky "Dead Duck" challenge, which goes like this: "If all ducks can fly and if Charlie is a duck, then Charlie can fly--unless Charlie is dead, in which case he cannot fly."

This is the type of thing computers don't infer, or at least, readily infer.

We can now consider another fundamental question-- how important, after all, is reasoning ability to intelligence?

Herbert Simon and his colleagues have conducted experiments on the way experts and novices solve physics problems.

In these experiments, it was soon determined that novices go about solving problems in a way similar to that of AI programs, i.e., "they work out elaborate strategies, complete with goals and subgoals, and solve every equation in sight. At each step, they ask "What do I do next?!".

"Experts, on the other hand, seem to look at a problem and see it whole: the solution is instantaneous. Whatever is going on, it is less a matter of reasoning than of recognition, of pattern matching on a huge store of memory and experience. Appropriate problem solving sequences are simply there when they are needed. It seems that knowledge, not reasoning, is power."

Researchers have so little understanding of how the brain works that they are of little help in AI work. It has been said that the inner workings of the brain are less familiar to us than the bottom of the sea or the geology of the moon.

According to J. Alan Robinson, head of AI research at Stanford University, the usual approach of AI workers is to ask themselves "If you were God, how would you have done it?"

Jacob Schwartz of NYU concurs. "The questions are deep and baffling," he states, "The human mind is doing very mysterious things!"

Russell Kirsch, head of AI research at the National Bureau of Standards says that, so far, "our little stunts are pretty weak soup."

Intelligence truly is a moving target.

Papert has the final word: "We are to thinking as the Victorians

were to sex. Everybody does it, but no one knows how to talk about it."

Concerning computer mechanics with the present state of the art, there are surprises for the outsider. AI computer schemes depart from standard expectations, for they don't crunch numbers; essentially, they manipulate symbols. They don't follow a rigid, precisely defined algorithm, instead they pick their own way through a problem according to the store of data, facts, and heuristic rules of thumb about the world.

Insights derived from AI's first quarter century are as follows:

1. Machines can behave intelligently using just two basic ingredients, search and knowledge.

2. None of the existing expert systems can learn in any real sense. (There is currently an effort to develop programs that can learn.)

3. In all applications there is to date a complete lack of anything which might be termed common sense. We humans stress common sense as an indispensable ingredient of intelligence, and AI's dream is to write a program that can learn from experience. It may be a long time in coming. One offhand estimate by McCarthy is 300 years.

It now appears that LISP (List Processing Language) originated and developed by John McCarthy for AI programming, may in time supplant other languages, all the way from mainframes down to personal computers.

Parts of this Lisp Language revolutionized programming languages. McCarthy conceived the "IF-THEN" expression, now a part of all computer languages. The use of this instruction permits long

tree-like formations of speculation, by means of which the program can proceed on an appropriate line of thought-- after examining by numerous trials and errors the particular problem domain.

The Lisp language is described as having "a purity and mechanical grace unlike any other....an abstract beauty..There is..a cumbrous grace in the way a program moves forward in Lisp. The parentheses coil and uncoil, shell to shell to shell descending through the dolls (within dolls) of a problem, and step to step to step coming back up again."

While other programs used in commercial computing are designed essentially for mathematical calculations, Lisp's strength is in manipulating non-numerical symbols, the symbols in which we humans think,.. such as words, phrases or geometrical figures. These symbols are the medium of formal logic, which philosophers consider the purest form of thought-- and most AI projects are based on logic.

Symbolic processing takes electronic processing beyond the domain of quantities into the domain of qualities.

(LISP does have its detractors. My son, a computer programmer at TI, refers to the acronym as standing for "Lots of Insane Silly Parentheses." His company doesn't agree, They are making several hundred LISP machines for MIT).

The pace of development in AI in robotics, vision, language skills, problem solving, etc., has been reasonable, perhaps somewhat relaxed now for many years. Unfortunately, we must step up the pace, for we now have international competition. On April 14, 1982, Tokyo's MITI (Ministry of International Trade and Industry) announced its "Fifth Generation" computer project. The timing

was ten years, the funding \$1 billion, the goal: computers three orders of magnitude faster than today's machines; expert systems capable of tapping knowledge bases as huge and diverse as Encyclopedia Britannica ; a natural language system which could translate from Japanese to English and back; and a high level interface with the ability to read, recognize images, and talk to human operators just as another human could do.

With the Japanese openly challenging us, the U.S. and others moved quickly. Britain began its Alvey Programme; on the continent there were attempts to begin an international program to be called Esprit. The latest entry in the fifth generation race is the Soviet Union. Their progress may depend largely on their skills in espionage.

In the U.S. 13 companies formed a research consortium under the acronym MCC (Microelectronics and Computer Technology Corporation). The plan is to spend \$1 billion in 10 years on AI, advanced microelectronics, and advanced computer architectures . In addition the Pentagon, under the acronym DARPA (Defense Advanced Research Projects Agency) has its funding for its own 10 year \$ 1-billion "Strategic Computing" program.

Considering the source of the funding one should expect primarily military and practical business applications. Basic work on AI may be slighted.

Roger Schank, of Yale University, tells us flatly " AI is in chaos. It's hard to get good researchers to work on the fundamental problems because the companies are snapping them all up. Theory has stagnated for the moment, and we've lost our momentum."

Among insiders, those best in position to judge, there is a

consensus that the public's expectations for AI have become dangerously overheated, that there is a real danger of severe disappointment. To the experts there is a sense that the whole MCC/DARPA/Fifth generation assault is just a self-exciting system, fed by the fuel of media hype. These researchers know that nothing in the science of AI has changed in the last 5 years. So-- why should we expect to make a quantum leap in the next ten years?

Roger Schank, Chairman of Yale's Department of Computer Science and Psychology, and director of the Yale Artificial Intelligence Program, at the age of 37, is interested in the nature of the mind. He says "Most of my work is in the structure of knowledge-- in learning and memory. I want to know how you know what you know."

He insists that the best way to build an intelligent machine is to imitate the human conceptual mechanisms that deal with language. For more than fifteen years he has worked toward the goal of getting computers to understand English, or just to get computers to understand. Eng

In his opinion, we cannot discuss language without simultaneously discussing memory, for these two are intertwined. This assumption led to the script concept. Schank explains it: "A script is a set of expectations, a codified set of information that seems to be associated in the mind with a particular event and that allows the inference process to be constrained. Suppose I tell you that I went to a restaurant and ordered lobster and that I paid the check and left. What did I eat? Well, I didn't say anything about eating, but it must have been lobster. Did the management get any money out of it? Of course, although I didn't say anything about management or money. Did the waitress give good service? What

waitress?"

"Language has a way of introducing lots of concepts sort of subliminally. When I talk about restaurants, I bring into your mind all the knowledge you have about ordering and waitresses and menus and tipping. A restaurant script.... the idea of scripts allows the computer to make more intelligent inferences because it knows about the domain in which it is operating."

Schank attacks the idea of formal logic " The classic view... is that the best way to build an intelligent system is to base it on formal logic. There are all sorts of people working on formal systems, but I was never comfortable with them. They ignore memory and learning and expectations because these are basically fuzzy, informal ideas. And I kept getting frustrated at how badly the formal systems worked.... people don't answer questions by proving theorems. For me it becomes a sort of existence proof. People do it and therefore I can get a machine to model it. It's perfectly plausible that there is another way to do it, but all my research says there isn't."

From these comments it can be seen that AI is yet in the early stages of attempting to build thought on an electronic as opposed to a neural substrate. The idea, until now, is simply to skim off the "symbol" level of the brain, or mind, and to duplicate the intermediate working levels of the brain. Oddly enough, what humans consider most difficult, mathematics, can be readily skimmed and adapted to computer work, while such a lowly virtue as common, or horse sense, completely eludes the programmer's attack.

As improvements are made, we continually revise our ideas about what constitutes intelligence. We are, in the AI approach,

essentially finding out what intelligence is not, to the possible benefit of psychologists.

A central question not yet answered is this: Must we completely duplicate the human brain to achieve the as yet undefined ultimate goals of AI? Skimming and mechanical tricks in binary code, which theoretically should work, may not work at all.

To achieve a robot comparable to the human frame and capability is, as yet, a wild and distant dream. As an example of our present limitations, a small robotic device with artificial vision capabilities managed recently to traverse a long, cluttered corridor without bumping into anything. To do so, however, it had to stop every three feet, scan and compute for fifteen minutes, then go forward another three feet. (As LBJ put it, it couldn't walk and chew gum at the same time). Obviously, speed is needed.

To get such speed, researchers now are trying parallel hookups, and banks of computers, in one case a setup to contain a million microchips. One shudders at the overwhelming problems in programming such arrangements.

Researchers in essence are nibbling away at the very fringes of the AI problem. The goals are nebulous. In fact, the ultimate artificial intelligence may not resemble human intelligence any more than jet planes resemble birds, their own organic counterparts. There is no way, at this stage, to project trend lines and see where all this will lead.

Of one thing we may be sure. In the process of getting there, we will certainly learn much more about ourselves.

Now, in conclusion, let's be optimistic, say it can be done, and ponder the implications of artificial intelligence.

Until now, all life and all intelligence have been carbon based. We humans are marvels of organic chemistry, carbon based. This sooty element combines easily with other elements to form so complex compounds, the building materials of life as we know it.

Another of our elements is silicon. Can it be that silicon is now on its way to life? Its path to life would be through the electronics industry, which even now is the computer's reproductive system. One might say that humans are the computer's way of making more computers, just as, to use Samuel Butler's famous ~~arg~~ argument, a chicken is simply an egg's way of making another egg.

While there is no real agreement as to the time frame, some have estimated that in about 20 years we will have the ability to construct a machine every bit as complex as our own ~~th~~ three pound soggy human computer. Add 50 years, and these machines could be our equals. From that point, with the built-in ability to design and improve themselves, they will reproduce and surpass us. This may prove inevitable. Carbon's organic monopoly on sentient life will be a matter of history: the lid of Pandora's box will be wide open.

Into the world will fly such problems as robotic rights, robotic freedom from discrimination, the robotic right to reproduce, perhaps even the right to robotic welfare (free electricity and batteries) medicare ( chip transplants and rust inhibitors). And, does the possession of highly advanced mental abilities imply the possession of a soul? Many other problems of a very serious nature could arise: would man survive? If so, in what relation to the superior intellects?

The human mind conceivably could become obsolete.

Perhaps, one hundred years hence, here in Hopkinsville, forty intellectuals will meet on first Thursdays, sit down, charge their batteries, and present two new data discs. These intellectuals will be members of The New and Improved Athenaeum Silicon and Robotic Society.