ARTIFICIAL INTELLIGENCE

by Harvey B. Sarles

Are humans machines? Can machines be like humans?

Such questions were posed first several centuries ago, when thinkers like Descartes and La Mettrie were faced with mechanical clocks and other devices which seemed-on their face-to do things which were, formerly, only within the domain of humans. [1]

To the first question, Hobbes and others began, in the 17th century, to use mechanical metaphors for thinking and intelligence. As well, the exploration and explanation of the body-as-machine, the mechanistic approach, emerged as paradigmatic; ancient puzzles such as the circulation of the blood, were solved. Anatomy and physiology, at least, used the model of the body-mechanic, to tell us a great deal about the nature of being human.

Currently, the quick retort to the first question-are humans machines-is "yes" and "no." The essential human, the "anima" of Aristotle, the "soul" or "spirit" of many others, is being pushed about by tissue transplants (even across species in the case of baboon heart being implanted in a human infant), and by such techniques as artificial insemination. In the context of the kinds of dualism which have characterized Western thought, the body-mechanic had been given short schrift. Now, it has raised serious questions of human and personal identity in our thinking. The importance and impact of the question has moved from the distant and the theoretical to the immediate and personal. Who anyone is (Who I am), is no longer as clear as it seemed to be, not too many years ago.

Are machines human-can they be?-is answered in short: "no" and "yes." No, they are not human; they lack "intelligence." The metaphors of Pinocchio and the Robot now entertain us for real. We have witnessed the rise of the age of computers, and the winds are shifting from an "intellectual" or "scientific" approach to the problem of "intelligence" to a technological or engineering ("expert systems") solution, in which very large computational systems will be made to imitate our best human thinkers. We stand on the verge of an extraordinary irony: computers will become "more" human, more intellectual through an approach to intelligence which is itself, mostly non-intellectual. The dystopic vision of the movie "2001" -in which thinking computers first rid the world of humanoids, then mimic us at our most destructive-impinges uncomfortably.

But before approaching a description of what is (as far as it can be foretold) happening in the field(s) called "Artificial Intelligence," it is useful to explore why these questions arise in the form that they do; i.e., as "intelligence." Quite probably, the human-machine problem could be considered more broadly, as an aspect of what makes humans, Human. If we

open the pages of our history, we see, however, that there are three significant counter-or anti-metaphors in terms of how we think about what is human, which equate human with intelligence. These are: animals, machines, and aliens (non-terrestrial beings, "Martians"). The question of intelligence and machines gains its substance largely in terms not of what we are; but of what we are not! As we will see, this historical habit of thought has depicted humans narrowly and less complexly, than we are...apparently.

We are not animals! Plato, who set the problems and the paths for their solution in Western thought, attempted to define humans in terms of some notion of "pure reason." It was cast within what I will call the "problem of knowledge" (epistemology), how to reconcile two aspects of knowing: the finite and the infinite. The body was held to be finite; the mind, to be infinite. Futher , he thought that we humans have-in common with animals -the finite body: we are thus both like animals ("in" nature); and different (outside of nature).

What is human, then, Plato defined as what is "unique" and infinite about humans...compared or juxtaposed with animals. This was satisfied, it turns out by "language." Humans, alone have "it"; other species do not.

Most of the history of thought, most of what we mean by the notion of "intelligence," is thus an examination, an analysis, or a critique of language: not how we do or behave, but how we express thought. Included in this have been: logic, syntax, words, and particularly, meaning; in general, how and what we know. In modern parlance, this has been a "mapping of the mind." [2]

And, it turns out also, it has been primarily a mapping of the individual mind: a kind of "psychology" of mind; leaving out or neglecting problems such as how we know or understand other persons; offering as solution to knowledge, certain theisms and texts which seemed, at least, to resolve the puzzle of human infinitude by attributing omniscience to a deity.

Without complicating this beyond anyone's comprehension, it can be repeated that the problem of what is Artificial Intelligence (AI), comes down to us pretty much as the attempt to analyze language. As we shall see, the MIT, or East Coast approach to AI, has been the analysis of language pretty much as it was laid out by Ludwig Wittgenstein early this century.' Further, we will see that this analysis of intelligence appears much more complex than it had earlier. Language, especially the ancient problems of meaning and context, includes much more than the organization of linguistic structures which is the study of syntax and grammar. The concept of language as being unique to humans, and consisting of grammer and syntax, was apparently oversimplified because we had compared humans with a vision of animals which was, itself, much too simple.

Second, humans are not machines. Machines were seen by Descartes, as purely mechanical. Like animals, they lacked a soul, a mind. Heirs to this Cartesian analysis of machines, we assumed that they could not think. The famous Cartesian "solution" set the problem of existence: the *"cogito ergo sum,* "-1 think, therefore I am. Machines could not think, essentially, by definition.

The current conceptual difficulty, is that machines seem to be able to do more and more which overlaps with what we have considered to be thinking. This ranges from being able to play games (e.g., chess), with a fixed and finite (if very large) set of rules, playernd domain, to being able to handle symbolic notions, to rewrite themselves (to "learn?"). This poses a dilemma, and places the original questions concerning humans and machines in a now blurred matrix: machines surely seem more human than they used to (they can do "intelligent" things, formerly reserved only to humans); we seem more like machines, reraising the question of what is human, in new and problematic ways.

Humans are not Martians...or maybe we are! The concept of communicating with extra-terrestrial beings-perhaps it is a sub-topic of humans as not-animals-has excited many thinkers, especially Carl Sagan. [4] Essentially, it presumes like Descartes, that only humans can think and have intelligence, at least upon this planet. But, other planets are likely also to possess "intelligent" life (i.e., like humans). To probe intelligence, conceived usually in some fairly pure form of rational thought as mathematics, we cannot compare ourselves with other terrestrial species (animals), but have to leave the earth. The importance of this as metaphor, is that it seems to "purify" the notion of intelligence, in the direction of formal or mathematical thought, and places the problem of AI primarily in the hands of mathematicians.

So the problem of AI has come down to us as the problem of what is (human) intelligence. Its focus has been the analysis of language "because" humans uniquely possess language; because language is the domain where we are (our minds are) infinite or indefinite in scope-responding to the problem of knowledge where humans are conceived to be both finite and infinite. Our finiteness has been compared to animals and to machines; our infiniteness is our intelligence, our language. How, then, to analyse language?

At this point, I would like to state the current "state of the art" in AI -much of which promises to "by-pass" these questions of what is intelligence, what is language-before returning to the question of how to analyze language, the approach which ponders, intellectually, what is intelligence.

Instead, there is a practical way around these questions, which leads, nonetheless, to "intelligent machines," by almost every definition. Enter the "engineering" solution! On the West Coast, at Stanford, the approach to AI is radically different from the East Coast, MIT, scientific-intellectual

approach. The Stanford group proposes not to probe what is intelligence, but to mimic what we know, certainly, to be intelligent. This is the "expert systems, knowledge-based" systems approach.

Backed by industry and the military, talking in 1986 to many thousands of engineers and industry managers by two-way teleconferencing means, Texas Instruments has presented Stanford's Ed Feigenbaum as the prophet and guru of this movement.'

What is proposed is simply to copy, imitate, or model the thinking and knowledge of the people who are already experts in a wide variety of fields: already suggested are physicians (the proverbial world's greatest surgeon), trouble-shooters who fix the machines and maintenance equipment of manufacturers, the person who places the just-right amounts of "xyz" in chemical and other stews, etc., etc. In the present, fairly excited climate, the question of irifinite human intelligence has been replaced by the limitlessness of the human imagination to define AI.

This entails the ability to imitate symbolically on very large computer, the ways an expert goes about seeing and thinking, gathering information, looking for more, ruling-out, rethinking, relearning...in short, getting "inside the heads" of those people who know best. If the proverbial country doctor wants to know how famous surgeon Dr. Q. thinks (or thought: he/she might now be dead), then the country doctor enters the program of Dr. Q with certain symptoms, observations, history, and will probe how Dr. Q. would go about diagnosing and treating such a patient. If a chemical stew looks or smells or reacts in some particular way, how would, say 3-M's expert, have gone about adding or correcting or aborting some process or procedure? The expert system becomes a kind of human manual in which the very best thinker is a dynamic step-by-step guide to today's doings.

This requires very large computers: and we are on the eve of enlarging vastly the capacities of computers by shifting from serial to parallel computing. Serial computing runs serially through a program, each step following the last. Parallel computers (e.g., many small computers linked by a master program which farms out work, or micro-chips with specialized functions), are within touch, and will be able to handle the 5,000 or so ways of considering thiT'./!,sthat experts are reputed to be able to do, in close to "real time."

The tough part, so far, is not in stating to the computer program how to imitate thinking-in fact, symbolic methods already appear adequate-but in getting persons who know how to interview in depth, the experts: who can get to know them, or get the experts themselves to become good self-analysts, to state how they know and do what they know and do. (Some old-fashioned "field-linguists" used to do approximately this type of work. It requires great patience, and ways of "getting into" another person's head: so-called "cognitive-mapping. ") A model for this approach, already worked out, is the "programmed" self-teaching course, some of which are quite effective.

In short, we are on the brink of a powerful, perhaps revolutionary shift in AI, toward the engineering approach. It is "non-intellectual" in the sense that it does not ask deeply about the nature of intelligence or language or human thought. It accepts that experts are, in fact, knowledgeable, and bypasses questions about the nature of that knowledge. It simply models their knowledge about the practical affairs of the world. It responds to the question of whether machines can think in a way that is as surprising as it is powerful: it makes no difference whether machines can actually think, as long as they can model human thought. The question of whether machines can think is, at once, moot, and answered in the affirmative!

Now, heading back East to MIT (a group in Pittsburgh at Carnegie-Mellon are no slouches either), the question of what is intelligent remains of great interest. It is not overly strong to state that applied math, philosophy, psychology, and linguistics (and ???), have joined forces to make computers be able to model what is intelligence. AI, in this context, has been placed within the development of powerful computer programs, which are increasingly geared to act intelligently. Here, the history of computing machines from Babbage to Turing to von Neumann, from simple recursion programs to complex branching and subroutine systems like Fortran and Pascal, to the vastly complex LISP programs, have increased the power of computing beyond the belief and comprehension of most of us. Here, AI is principally the ingenuity of thinkers, mostly mathematicians, to concoct programs which act much like humans: which "think."

If they "solve" the question of what is intelligence, and can program "it," then the computer would presumably be able to think or learn, change its "mind" (its program), grow in scope from a manipulator of information, toward becoming knowledgeable, and, with great difficulty (my present perception), to become wise.

It appears, from my vantage, that this intellectual-language analysis field is presently at some impasse. The (early) Wittgensteinian approach to language was too "simple," and didn't account of semantics of meaning. Why it was too simple has to do with the history of philosophy underestimating the cleverness of other animals, thence of humans.

But it-language grammar, syntax, structure as presently conceived-is not very capable of dealing with many questions which seem very simple for human beings: e.g., context, what is implied but not stated, mutual understanding, intelligibility. Much of the current and forseeable progress in AI, East Coast led, would seem to be directed toward deepening understanding of human cognitive processes. Interesting, ambitious; but not close to "solving" or resolving, practically, the ability of AI to act intelligently, as the Stanford, engineering approach promises to do, imminently. Can machines think: here, not yet, not for quite a while.

Other aspects of the problem of AI, conceived scientifically or intellectually (I tend to equate these), is that "intelligence" can be thought about in many more ways, than the philosophical-mathematical approach includes.

Probably, this is also true of how humans think about intelligence; or think intelligently. There are many "mappings" or maps of the mind: familiar ones include the Freudian; perhaps the sign language of the deaf; the world-views of other cultures which have different visions of humans (and other species) including Confucian, Amerindian, South Asian concepts of what is being and thinking and existence; transcendental theories including Emerson and Thoreau and Nietzsche; Anthropological theories which concentrate first upon social interaction, talk is "for" communicating and understanding (my own work on an interactional Question-Response Grammar); etc., etc.6,7

Much more is occurring in the field which is called AI, than has been indicated here: the development of aspects of AI such as machines which we can talk into directly, addressing computers and all their outputs; machines designed especially for the handicapped, permitting the blind to "see," the deaf to "speak," the paralyzed to move effectively and with ease.8

However, in broad outline, the foreseeable future seems to be moving, inexorably and with great power (rivaling in importance, perhaps, the potential effects of genetic engineering), toward expert and knowledge- based systems to make machines act as if they are intelligent. The question of whether machines are realy intelligent will thus be by-passed: If they act as if they are intelligent, who can say that they are not?

Expert systems also address, the question of whether humans are machines in a peculiar way: it will be increasingly difficult, as we model the best and most knowledgeable humans, to distinguish between the thinking and output of persons and machines. Whether humans are machines, will be posed in the context of significantly smaller or different distinctions than previously.

The important human questions remain, still largely unexamined, in the current exuberance of "can-do." How are we to deal with the machines and with ourselves: do we become "more human," enlarged in scope and outlook, morally activated, using the leisure and possibilities of these machines toward human ends; or will we become diminished and hand-maidens to our own devices, dehumanized and relying upon technology to tell ourselves who we are?

REFERENCES CITED

Keith Gunderson. *Mentality and Machines.* Anchor Books. Garden City, N.Y.: Doubleday & Co. 1971.

Charles Hampden-Turner. *Maps of the Mind.* New York: Macmillan, 1981.

Raymond Kurzweil. "Whate is Artificial Intelligence Anyway? *American Scientist,* Vol. 73, May-June 1895.

Carl Sagan, Ed. *Communication with Extraterrestrial Intelligence.* CETI. Cambridge: MIT Press 1974.

"Knowledge-Based Systems: A Step-by-Step Guide to Getting Started." *The Second Artificial Intelligence Satellite Symposium.* Symposium Notes. Texas Instruments, Sponsor.

Harvey B. Sarles, *Language and Human Nature.* Minneapolis: University of Minnesota Press, 1985.

Robert J. Sternberg: "Human Intelligence: The Model is the Message." *Science* 230 (4730), 6 Dec 1985.

Kurzweil, op cit.