AI Magazine Volume 5 Number 4 (1984) (© AAAI)

Intelligence, Artificial and Otherwise

William M. Chace
Stanford University, April 27, 1984

I RISE NOW TO SPEAK with the assumption that all of you know very well what I am going to say. I am the humanist here, the professor of English. We humanists, when asked to speak on questions of science and technology, are notorious for offering an embarrassed and ignorant respect toward those matters, a respect, however, which can all too quickly degenerate into insolent condescension. Face to face with the reality of computer technology, say, or with "artificial intelligence," we humanists are the kind of soreheads who compulsively point out that human beings aren't machines, that computers will never possess the uniquely human powers of intuition and common sense, and that we smelly, hairy, and otherwise organic *people* are simply more interesting, more clever and more mysterious than electric-powered, binary-formatted, digital computers with their tidy little green phosphorescent screens.

So now you know what I'm going to say, or at least part of it. For, to be even more candid with you this morning, I want to go on to say that while I do believe many of the things that we humanists are reputed to believe, I also happen to believe that humanists are wrong, and very foolish, to fear computers. But they are also wrong to underestimate the power of those computers and their important utility in the world that is now ours and the world that we will soon be turning over to our descendents.

Many of my colleagues at Stanford and elsewhere, and I will presume many of you, feel intimidated or very inept in the presence of computing power. We have become anxious in the presence of the dazzling competence, the accuracy, the speed, and the enormous store of memory that characterizes the digital computer. But I come before you to say that our anxieties are baseless. At the same time, our ignorance is great. Computing power, as I have come to understand it as a rank amateur user of an IBM-PC, is not what I had once thought it to be. It is both stronger and more capable than I had believed and, at the same time, more limited and less awesome. I am, as a humanist, on my way to demystifying some of the powers of the computer, at least to my own satisfaction, and I invite you for the next few minutes or so to share with me the grounds for this demystification.

But first we must attend to the awkwardness and the suspicion we feel. Why are we so anxious? The reasons, I think, are historical. In the last 125 years, the human community has suffered at least three traumatic shocks to its sense of well-being and primacy in the universe. These shocks have been the Darwinian, the Marxian, and the Freudian.

Darwin quietly and patiently assembled the evidence to revel the individual human being's true place with respect to his phylogenetic antecedents and his phylogenetic posterity. After Darwin's biological work was done, his followers and his publicists claimed that individual human units (us, for instance) were, in sum, trivial and ephemeral; biological and genetic evolution was, in comparison, all-important. The result of the Darwinian revolution, for dignity and the self-regard of the particular human being (you, for instance), was crushing. And it was particularly crushing to that individual's sense of divine inspiration and "connectedness." Organized religion was placed on the defensive, for mankind's allegedly "special" relationship to God had been radically questioned.

The writings of Karl Marx also had the effect of weakening the human sense of self-importance. Marx attacked the ideal of the culturally "free" and independent thinker and thought "pure reason" an unattainable goal for humans. He claimed that what you consciously *think* you are, and what you think you are doing, in the world surrounding you, is thinking wholly conditioned and shaped by the social and economic conditions of that surrounding world. You can, said Marx, never lay full claim to the "truth" about your life in society. That is because your consciousness, molded by social, historical, and economic forces, always obscures such truth. In 1859, the same year as *The Origin of Species*, Marx wrote that "it is not the consciousness of men that determines their social being, but their social being that determines their consciousness." With such declarations, he undercut the belief in the autonomy of the solitary individual and the notion that such an individual could exercise "free choice."

Sigmund Freud further deflated the importance and the value of human consciousness, for he proclaimed that con-

---

Some 1,200 Stanford alumni gathered in New York City April 26-29, 1984 for "Life in the 21st Century," the largest alumni program ever sponsored by the University One of the highlights was an exchange between Prof. Edward Feigenbaum, computer science department, and Prof William Chace, English department, associate dean of Humanities and Sciences Following are Chace's prepared comments on "Intelligence, Artificial and Otherwise", reprinted with permision from *Campus Report* (Stanford University: 9 May 1984), 17-18

sciousness was but a thin veneer atop a turbulent and opaque sea of repressed desires, anxieties and inexpressible conflicts left over from our individual pasts. Of course we could possess consciousness, he said, but that was to possess almost nothing. The true springs of human motivation and character formation were submerged beyond final recovery and, try as we might, we could never truly achieve the privilege and nobility, first celebrated by the ancient Greeks, of "self-knowledge." That kind of knowledge was, he argued, to be withheld permanently from us; the interfering power of buried childhood conflicts was simply too great.

If these three blows to human self-esteem have been sharp and unforgettable, they have now been followed by yet one more. Thus I return, as you knew I would, to the impact of the computer revolution. That phenomenon has forced all of us, disheveled and shell-shocked after the first three traumas, to inquire: "What indeed am I particularly good for? Acknowledging that I am inefficient and inaccurate at remembering, say, airline schedules and factory inventories, I ask: at what am I uniquely skilled, since I can no longer do the binomial theorem in my head? Why, in time, cannot my computer replacement be found?" For computers do make us all feel thick-witted at memory tasks, at statistics and other operations involving quantity, at budgets and banking, and, indeed, doing any arithmetical work more advanced than that we once learned in the fifth grade. What's more, computers are more numerous today than yesterday, this hour more numerous than last. In 1970, there were 65,000 computers in the United States; now there are more than five million. And the amount of computing power is doubling every two years. What's more, the miniaturization of that computing power means that today you can put on your desk a machine that used to require the space of many large rooms.

So what, indeed, are humans– outwitted and on their way inevitably to further embarrassment –good at? I arrive at last at my question for today. And should you anticipate the response you would want me to make, and should you say, in a cheerily helpful way, that airline schedules and factory inventories should be left to the computer and more important matters, such as *judgment*, *expertise*, *intuitive discernment* and *the wisdom of experience* be kept under the control of human beings, let me quote to you an article from the *The New York Times* one month ago. It reports on the activities of a company located just down the road from Stanford University that will, with the help of a phenomenon called "knowledge engineering," attempt "to perform such diverse tasks as evaluating casualty risks, making commercial credit decisions and controlling oil-well drilling, tasks (the article goes on to say) that are extremely difficult and often done well only by a relatively small number of experts."

And how does this "knowledge engineering" work? The knowledge engineer first interviews a set of people expert at solving problems in science, medicine, business and other endeavors to find out how they make the judgments that are the core of their expertise. The knowledge engineer next codifies that knowledge so the computers he controls can make similar decisions by emulating the great power of human reasoning, which is the power to infer, deduce and combine elements. The expert's wisdom is then reduced to a large set of interconnected generalized rules called the "knowledge base." Next, a separate computer program called an "inferential engine" is used to search that knowledge base and draw judgments when confronted with evidence from a particular case, much the way an expert applies past knowledge to a new problem. Thus, for example, a knowledge engineer equipped with the appropriate knowledge base, and provided with a list of a given medical's patient's symptoms, might diagnose a certain kind of emphysema in that person.

An alleged advantage of this method is that the expert, who is now housed in a computer, has become disembodied and immortal. As the youthful owner of this particular company puts it, "in every organization there is usually one person who is really good, who(m) everybody calls for advice. He is usually promoted, so that he does not use his expertise anymore. We are trying to protect that expertise if that person quits, dies or retires and to disseminate it to a lot of other people." So much, then, for the question of the wisdom of experience; so much, then, for human judgment being left to humans. Wisdom, we read in the *Times*, can be reproduced at will and judgment can be left to computer storage.

Lest any of us think that this recent example of "knowledge engineering" or what we can call applied artificial intelligence is all quite dubious and far-fetched, I urge our attendance on the fact that artificial intelligence, since its modern inception in the minds of such extraordinary scientists as Alan Turing, John von Neumann, Herbert Simon, John McCarthy and my fellow panelist Ed Feigenbaum, has been notified on innumerable occasions that computer programs cannot, simply *cannot*, do certain things and that artificial intelligence is living in a world of fancy dreams and science fiction to think it can. But, truth to tell, some of those things, such as developing a computer chess-machine that can win on a very advanced level of competitive play, have been achieved. So have programs which, at a very high level of sophistication, can infer chemical structure from mass spectrometry.

The fact of the matter seems to be that artificial intelligence responds with enormous energy to the challenge of being told that it simply *can't* do something. If it is told that a certain mental maneuver is *particularly and uniquely* human, and therefore undoable, then AI proposes that such a maneuver will be its next achievement. In his wonderful book, *Gödel, Escher, Bach—An Eternal Golden Braid*, Douglas Hofstadter wittily observes that " ..once some mental function is programmed, people soon cease to consider it as an essential ingredient of 'real thinking.' The ineluctable core of intelligence is always in that next thing which hasn't yet been programmed ... AI is whatever hasn't been done yet."

It is my belief, then, that artificial intelligence and, as they charmingly call themselves, the artificial intelligentsia, will probably be able to mimic many particular human intellectual techniques of problem-solving, even the ones I most adamantly now *know* they won't be able to do. The chess machine will get better, the language translators will improve beyond their present pitiable level of performance, some short poems and short stories of rather dubious but not entirely laughable quality will be composed and a great range of precisely designated scientific tasks involving the analysis of myriad data will be performed at incredible rates

of speed and accuracy.

Some things, however, will continue to prove very difficult indeed for AI, such as creating a machine that will effectively ride a bicycle (for that involves, as we all know, great motor balancing and coordination); or creating a machine that will be able to perform several intellectual tasks in parallel fashion. And I think that certain other kinds of intellectual tasks, such as judging the quality of a painting, or of a musical composition, are fully beyond the horizons of artificial intelligence. Nor do I believe that "knowledge engineering" will enter the courtrooms of our country in the foreseeable future. Those courtrooms are arenas in which *human* judgment at its highest ranges can sometimes be found. In sum, I do not believe it will be possible to announce, sometime in the 21st century, that human intelligence and artificial intelligence have at last become one, and that the machine has at last learned to mimic, to replicate, *all* the workings of the human brain. And why not? The reasons, I think, are several:

I believe, first of all, that there will prove to be little real ambition on the part of specialists working in the community of artificial intelligence to take up such a task in all its dimensions. As we look around the globe and witness its enormous population, one of the first features that must strike our attention is just how little we need more things that think like human beings. Human beings think like human beings and of human beings we have an overwhelming superabundance. So, while the intellectual appeal of creating replicas of human mind in its full range and complexity will be strong, the practical demands will be weak.

The second reason is that artificial intelligence performs its tasks best when it knows with absolute precision what those tasks are. It works discretely and directly, always with a menu of problems to be solved, answers to be registered, goals to be reached and missions to be executed. To understand the workings of a computer is to come to appreciate a world of the completely explicit, a world from which nuances have been removed, a world in which a tolerance for ambiguity is not an important operational feature. Human intelligence, by contrast, is only spasmodically explicit and direct and is notable for both the waywardness of its functioning and the ambivalence it brings even to its most serious responsibilities. Prince Hamlet (if you will permit me something so foreign as a literary reference), knowing he must kill the king, does so only after a long period of noble procrastination and eloquent malingering. A good little piece of knowledge engineering would have wiped out wretched Claudius in the first act, leaving Hamlet with an agenda of other princely tasks and us without a play to watch.

The third reason is that successful imitations must operate with a genuine knowledge of that which is being imitated. For artificial intelligence to achieve its final goal, for it to arrive at that stunning moment when its procedures will become absolutely identical with natural intelligence—when machine mind is indistinguishable from human mind—one of the toughest questions in the history of philosophical and psychological investigation will have at last been answered. That question is, of course, "what is mind?" Are we satisfied with the answer Plato gave to this question? Are we happy with Descartes' answer? Or Locke's? Or Kant's? Or Wittgenstein's? Freud's? Jung's? Has any philosopher or any psychologist described mind in such a manner that we know, truly know, its contents, and the reservoir of its possibilities? The answer is certainly and decisively No.

Without, then, a firm knowledge of what it is attempting to simulate, artificial intelligence will be able, sometimes astonishingly well, only to mimic certain discrete and well-defined features of human mind. Of course, much to our surprise or perhaps to our dismay, it will simulate more and more of those features as time goes on. And we might be led to think that the results of its operations, results so much like the results of human thought, will be sufficient proof that the internal operations of the machines are wholly similar to the internal operations of human thinking. But equivalent results will not necessarily mean equivalent processes. The external products, so impressively akin, will not necessarily betoken internal likenesses.

Why is this so? What will prevent the machine from acquiring a genuine comprehension of the contents of human mind? Again, Douglas Hofstadter can help us, for in his speculations he says, "unless (the AI program) had an amazingly faithful replica of a human body...it would probably have enormously different perspectives on what is important, what is interesting, *etc.*" This must certainly be true. For, as your bodies have been telling your minds for the last few hours this morning, those minds live inside those bodies, and those bodies tire, grow restless, must move, and have countless little desires that must be answered. I know, and you know I know, that your minds have not been wholly attentive to everything I have been saying. Only part of the reason for this inattentiveness has been the tediousness of my remarks. The fact of the matter is that human beings—cellular, organic, hungry, thirsty and given to sleep—are sometimes quite bad at being responsive thinking machines.

I add to this description of human beings the fact that we are mortal. We contain within us both the knowledge and the seeds of our own necessary disappearance. We age, doing so at the instruction of something deep within us. The poet William Butler Yeats summed up the relationship of his creative and energetic mind to the body in which it was encased by saying:

> ". . sick with desire
> And fastened to a dying animal
> It knows not what it is."

We are, and we know we are, dying animals, and I submit that that particular piece of knowledge, and the peculiar melancholy surrounding it, will not easily be introduced into any routine, or sub-routine, of a program in AI. Why not? Because while mortality, like mind itself, is a reality many human beings have pondered, and while the results of such pondering have been symphonies, novels, poems and the like, no one among us, no one in our cultural ancestry, has gotten mortality down right. Who is the experts' expert on death? What understanding of the human relationship to individual extinction shall we relay to the "knowledge engineer"? Who has perfectly grasped the awareness we have of ourselves, and of our own certain death, when someone we love passes out of the world? I have no such person to nominate to the community of experts in artificial intelligence, or rather I have many nominations to make, all of them wise people, but I can make no final judgment as to which of them knows the truth of this matter.

We are in bodies. We are mortal. We are also, as people, capable of extraordinary powers of self-consciousness. We can, without limit, think about our thinking, and then think about the thinking about that thought, and proceed in an infinitely self-reflexive pattern of thinking about thought. Of course such self-consciousness is at last useless. It produces no results save the strange, even weird, satisfaction that it can be done. That, of course, is my very point. Human beings are good at what we can term "meta-thought" or "meta-cognition" and machines aren't. Human beings, moreover, are only occasionally productive, only occasionally goal-oriented. Machines, however, are best at being productive and goal-oriented. Much of the time humans day-dream, waste time, spin their wheels and, in general, prove inefficient and prodigal of their abilities. We are still the people our discouraged parents said we were: slothful. Not so with computers.

Indeed, to construct a computer programmed to imitate the infinite possibilities of inefficiency, listlessness, muddle-headedness, and blind confusion of which we human beings have proved ourselves so capable is, I suppose, possible—given sufficient man-hours and governmental and industrial support– but I think unlikely in the extreme. Indeed, I would imagine that my fellow panelists would think that to build a program with the particular specifications of Witlessness and Knuckleheadedness would go against the very grain of their professional discipline. Just as perfect nonsense might be much more difficult to generate than sense, so I think the task of comprehending the full possibilities of human foolishness might be infinitely more difficult than the generation of the rules governing successful problem-solving. Let us, then, leave to the historian the understanding of our folly, and to the computer scientist the grasp of our competence on those occasions when we are competent.

And, as long as I am in the business of dividing up the intellectual labor of our time and of the 21st century, I might offer a few words on the venerable subject of the cultural split between the "two cultures"—the scientific and the humanistic. This famous split, thrust into public view by C. P. Snow in the 1950s and the fuel for many a debate and public speech ever since, is one of the great red herrings of our time. Sir Charles seemed to believe that the gulf separating the scientific sensibility from the humanistic sensibility could be bridged—by specifically what he never said– perhaps by hard work, lots of good will and much reading, perhaps by himself.

But I do not believe the bridge can so easily be built. The world of science, of which the energies and fascination of artificial intelligence is one small part, is a world virtually closed to everyone save highly-trained experts and their students, some of them at Stanford. With each new real development in scientific knowledge, the lay person is left further behind. With each passing day, science really does augment its own storehouse of knowledge, and it does so with the help of highly sophisticated vocabularies inaccessible to almost everyone else. Not to grasp that fact is to pretend that science is less than it is, and thus to vulgarize and patronize science.

Things are otherwise in the "humanities," so-called, and in the arts in general. Although I have colleagues who despair of the embarrassment my remarks might cause, I have come before you to claim that the arts—literature, music, painting, and so forth—really have very little that is "new" or inaccessible or technically arcane to say to you. Its techniques might be difficult to grasp, but its "message" rarely is. Art is nothing if not a reminder of those truths you already knew. It is a solace, or a pang of guilty awareness, or a delight. It makes you recognize, with dramatic clarity, the meaning of it all. It startles and arrests you, and then pleases you by confirming your deepest intuitions. Thus art's pleasure.

In being so radically different an enterprise from scientific investigation, art seeks no cozy liaison with anything else. Nor would our time be well spent were we to try to force that liaison. We should learn to tolerate, with benign sympathy, the different ways in which the human mind can work. And that tolerance should permit us to recognize that only a very small handful of experts really understand how computers work. Some of them are with you today. They will, I assume, continue their work of designing programs of ever greater sophistication and power. The aim of their enterprise will be to develop more complex programs and procedures so that the rest of us, in employing computers for our tasks, can afford to be no more conscious of the computer than we are of the pencil or the telephone. After all, more people today are dealing with computers in an intuitive and natural manner because programmers have written software that makes it less necessary for users to know what the machine is actually doing. And I call this ignorance on our part useful and welcome. The computer is, when all is said and done, a tool, no more than a tool, and it is *our* tool. We will use it after craftsmen have made it available to us. So it was once with the hammer, cannon and the knife; so once it was with the automobile, the radio and the airplane. Originally objects of awe, curiosity, and even terror—we have exerted our full authority over them all. So now it is with the digital computer. We will, certainly in the 21st century, claim it is one more man-made instrument that we have controlled and demystified.

I leave you with the simple assertion that Frankensteinian monsters exist in books, not in life. Science has yet to create very few autonomous forces with which we cannot reckon. Perhaps nuclear energy is one, but computing power and artificial intelligence are not. In sum, kind listeners, I find that human beings, in all their mystery and ineptitude, are not about to be replaced. In saying that and in reminding you that we are both mysterious and inept, I think I am leaving you with both good news and bad news.

AAAI RECOGNIZED

AAAI's convention publication entry in the American Society of Association Executives' annual Gold Circle Awards competition was singled out by the judges for recognition. AAAI's entry will be displayed at the ASAE Management Conference, 2-5 December, 1984, to be held at the Hyatt Regency Crystal City Hotel in Arlington, Virginia.