THE MIND OF THE MACHINE

within the life span of many of us, superintelligent computers—almost interchangeable with man—will bring the mixed blessing of a world without work

article By ARTHUR C. CLARKE

OURS IS THE CENTURY in which all man's ancient dreams—and not a few of his nightmares—appear to be coming true. The conquest of the air, the transmutation of matter, journeys to the Moon, even the elixir of life—one by one, the marvelous visions of the past are becoming reality. And among them, the one most fraught with promise and peril is the machine that can think.

In some form or other, the idea of artificial intelligence goes back at least 3000 years. Talos, the metal man who guarded the coast of Crete, however, was only a physical and not an intellectual giant; perhaps a better prototype of the thinking machine is the brazen head generally linked with the name of Friar Bacon, though the legend precedes him by some centuries. This head was able to answer any question given to it, relating to past, present or future; but, as is customary with oracles, there was no guarantee that the inquirer would be pleased with what he heard.

Over these tales there usually hangs the aura of doom or horror associated with such names as Prometheus, Faust and, above all, Frankenstein, though that unfortunate scientist's creation was not a mechanical one. Perhaps the finest work in this genre is that little classic of Ambrose Bierce's, Moxon's Master, which opens with the words: "Are you serious? Do you really believe that a machine thinks?"

It will not be universally accepted, but there is one very straightforward answer to this question. It can be maintained that every man is perfectly familiar with at least one thinking machine, because he has a late-type model sitting on his shoulders. For if the brain is not a machine, what is it?

Critics of this viewpoint (who are probably now in the minority) may argue that the brain is in some fundamental way different from any nonliving device. But even if this is true, it does not follow that its functions cannot be duplicated, or even surpassed, by a nonorganic machine. Airplanes fly better than birds, though they are built of very life different materials.

For obvious psychological reasons, there are people who will never accept the possibility of artificial intelligence, and would deny its existence even if they encountered it. As I write these words, there is a chess game in progress between computers in California and Moscow; both are playing so badly that there is clearly no human cheating on either side. Yet no one really doubts that eventually the world champion will be a computer; and when that happens, the die-hards will retort: "Oh, wellchess doesn't involve real thinking," and will point to various grand masters in evidence.

One can sympathize with this attitude, but to resent the concept of a rational machine is itself irrational. We no longer become upset because machines are stronger, or swifter, or more dexterous than human beings, though it took us several painful centuries to adapt to this state of affairs. How our outlook has changed is well shown by the ballad of John Henry; today, we should regard a man who challenged a steam hammer as merely crazy-not heroic. I doubt if contests between calculating prodigies and electronic computers will ever provide inspiration for future folk songs. But I'll be happy to donate the theme to Tom Lehrer.

It is, of course, the advent of the modern computer that has brought the subject of thinking machines out of the realm of fantasy into the forefront of scientific research. One could not have a plainer answer to the question that Bierce posed three quarters of a century ago than this quotation from MacGowan and Ordway's recent book, Intelligence in the Universe: "It can be asserted without reservation that a general-purpose digital computer can think in every sense of the word. This is true no matter what definition of thinking is specified; the only requirement is that the definition of thinking be explicit."

That last phrase is, of course, the joker, for there must be almost as many definitions of thinking as there are thinkers; in the ultimate analysis, they probably all boil down to "Thinking is what I do." One neat way of avoiding this problem is a famous test proposed by the British mathematician Alan Turing, even before the digital computer existed. Turing visualized a "conversation" over a teleprinter circuit with an unseen entity "X." If, after some hours of talk, one could not decide whether there was a man or a machine at the other end of the line, it would have to be admitted that X was thinking.

There have been several attempts to apply this test in restricted areas-say, in conversations about the weather. One clever program (DOCTOR) has even allowed a computer to conduct a psychiatric interview, with such success that 60 percent of the patients refused to believe afterward that they were not "conversing" with a flesh-and-blood psychiatrist. But as people talking about themselves can be kept going indefinitely with a modest supply of phrases like "You don't say!" or "And then what did you do?," this particular example only demonstrates that little intelligence is involved in most conversation. The old gibe that women enjoy knitting because it gives them something to think about while they're talking is merely a special case of a far wider law, ample proof of which may be obtained at any cocktail party.

For the Turing test to be applied properly, the conversation should not be restricted to a single narrow field but should be allowed to range over the whole arena of human affairs. ("Read any good books lately?" "Has your wife found out yet?," etc.) We are certainly nowhere near building a machine that can fool many of the people for much of the time; sooner or later, today's models give themselves away by irrelevant answers that show only too clearly that their replies are, indeed, "mechanical," and that they have no real understanding of what is going on. As Oliver Selfridge of MIT has remarked sourly: "Even among those who believe that computers can think, there are few these days, except for a rabid fringe, who hold that they actually are thinking."

Though this may be the generally



THE MIND OF THE MACHINE

within the life span of many of us, superintelligent computers—almost interchangeable with man—will bring the mixed blessing of a world without work

article By ARTHUR C. CLARKE

accepted position in the late 1960s, it is the "rabid fringe" who will be right in the long run. The current arguments about machine intelligence will slowly fade out, as it becomes less and less possible to draw a line between human and electronic achievements. To quote another MIT scientist-Marvin Minsky, professor of electrical engineering: "As the machine improves . . . we shall begin to see all the phenomena associated with the terms 'consciousness,' 'intuition' and 'intelligence' itself. It is hard to say how close we are to this threshold, but once it is crossed, the world will not be the same.... It is unreasonable to think that machines could become nearly as intelligent as we are and then stop, or to suppose that we will always be able to compete with them in wit and wisdom. Whether or not we could retain some sort of control of the machines, assuming that we would want to, the nature of our activities and aspirations would be changed utterly by the presence on earth of intellectually superior beings."

Very few, if any, studies of the social impact of computers have yet faced up to the problems posed by this last sentence-particularly the ominous phrase "assuming that we would want to." This is understandable; the electronic revolution has been so swift that those involved in it have barely had time to think about the present, let alone the day after tomorrow. Moreover, the fact that today's computers are very obviously not "intellectually superior" has given a false sense of security-like that felt by the 1900 buggy-whip manufacturer every time he saw a broken-down automobile by the wayside. This comfortable illusion is fostered by the endless stories-part of the transient folklore of our age-about stupid computers that have had to be replaced by good oldfashioned human beings, after they had insisted on sending out bills for \$1,000,000,004.95, or threatening legal action if outstanding debts of \$0.00 were not settled immediately. The fact that these gaffes are almost invariably due to oversights by human programers is seldom mentioned.

Though we have to live and work

with (and against) today's mechanical morons, their deficiencies should not blind us to the future. In particular, it should be realized that as soon as the borders of electronic intelligence are passed, there will be a kind of chain reaction, because the machines will rapidly improve themselves. In a very few generations—*computer* generations, which by this time may last only a few months—there will be a mental explosion; the merely intelligent machine will swiftly give way to the *ultra*intelligent machine.

One scientist who has given much thought to this matter is Dr. Irving John Good, of Trinity College, Oxford-author of papers with such challenging titles as "Can an Android Feel Pain?" (This term for artificial man, incidentally, is older than generally believed. I had always assumed that it was a product of the modern science-fiction magazines, and was astonished to come across "The Brazen Android" in an Atlantic Monthly for 1891.) Dr. Good has written: "If we build an ultraintelligent machine, we will be playing with fire. We have played with fire before, and it helped keep the other animals at bay."

Well, yes---but when the ultraintelligent machine arrives, we may be the "other animals"; and look what's happened to them.

It is Dr. Good's belief that the very survival of our civilization may depend upon the building of such instrumentalities; because if they are, indeed, more intelligent than we are, they can answer all our questions and solve all our problems. As he puts it in one elegiac phrase: "The first ultraintelligent machine is the last invention that man need make."

Need is the operative word here. Perhaps 99 percent of all the men who have ever lived have known only need; they have been driven by necessity and have not been allowed the luxury of choice. In the future, this will no longer be true. It may be the greatest virtue of the ultraintelligent machine that it will force us to think about the purpose and meaning of human existence. It will compel us to make some far-reaching and perhaps painful decisions, just as thermonuclear weapons have made us face the realities of war and aggression, after 5000 years of pious jabber.

These long-range philosophical implications of machine intelligence obviously far transcend today's more immediate worries about automation and unemployment. Somewhat ironically, these fears are both well grounded and premature. Although automation has already been blamed for the loss of many jobs, the evidence indicates that so far, it has created many more opportunities for work than it has destroyed. (True, this is small consolation for the particular semiskilled worker who has just been replaced by a couple of milligrams of microelectronics.) Fortune magazine, in a hopeful attempt at selffulfilling prophecy, has declaimed: "The computer will doubtless go down in history not as the explosion that blew unemployment through the roof but as the technological triumph that enabled the U.S. economy to maintain the secular growth on which its greatness depends." I suspect that this statement may be true for some decades to come; but I also suspect that historians (human and otherwise) of the late 21st Century would regard that "doubtless" with wry amusement.

For the plain fact is that long before that date, the talents and capabilities of the average-and even the superiorman will be as unsalable in the market place as his muscle power. Only a few specialized and distinctly non-white-collar jobs will remain the prerogative of nonmechanical labor; one cannot easily picture a robot handy man, gardener, construction worker, fisherman. . . . These are professions that require mobility, dexterity, alertness and general adaptability-for no two tasks are precisely the same-but not a high degree of intelligence or data-processing power. And even these relatively few occupations will probably be invaded by a rival and frequently superior labor force from the animal kingdom; for one of the longrange technological benefits of the space. program (though no one has said much about it yet, (continued on page 122)



MIND MACHINE (continued from page 119) Þ

for fear of upsetting the trade unions) 60 will be a supply of educable anthropoids filling the gap between man and the be great apes. 1

0

H

It must be clearly understood, therefore, that the main problem of the future-and a future that may be witnessed by many who are alive todaywill be the construction of social systems based on the principle not of full employment but rather full unemployment. Some writers have suggested that the only way to solve this problem is to pay people to be consumers; Fred Pohl, in his amusing short story The Midas Plague, described a society in which you would be in real trouble unless you used up your full quota of goods poured out by the automatic factories. If this proves to be the pattern of the future, then today's welfare states represent only the most feeble and faltering steps toward it. The recent uproar about Medicare will seem completely incomprehensible to a generation that assumes every man's right to a basic income of \$1000 a year, starting at birth. (In New Dollars, of course; 1 N. D.=\$100, 1984 currency.)

I leave others to work out the practical details of an economic (if that is the right name for it) system in which it was antisocial, and possibly illegal, not to wear out a suit every week, or to eat three six-course meals a day, or to throw away last month's car. Though I do not take this picture very seriously, it should serve as a reminder that tomorrow's world may differ from ours so radically that such terms as labor, capital, communism, private enterprise, state control will have changed their meanings completely-if, indeed, they are still in use. At the very least, we may expect a society that no longer regards work as meritorious or leisure as one of the Devil's more ingenious devices. Even today, there is not much left of the old puritan ethic; automation will drive the last nails into its coffin.

The need for such a change of outlook has been well put by the British science writer Nigel Calder in his remarkable book The Environment Game; "Work was an invention, which can be dated to the invention of agriculture. . . . Now, with the beginning of automation, we have to anticipate a time when we must disinvent work and rid our minds of the inculcated habit."

The disinvention of work: What would Horatio Alger have thought of that concept? Calder's thesis (too complex to do more than summarize here) is that man is now coming to the end of his brief 10,000-year agricultural episode; for a period of a hundred times longer he was a hunter, and any hunter will indignantly deny that his occupation 122 is "work." We now have to abandon

agriculture for more efficient technologies; first, because it has patently failed to feed the exploding population; second, because it has compelled 500 generations of men to live abnormal-in fact, artificial-lives of repetitive, boring toil. Hence, many of our present psychological problems; to quote Calder again: "If men were intended to work the soil, they would have longer arms."

"If men were intended to . . ." is, of course, a game that everyone can play; my favorite competitor is the old lady who objected to space exploration because we should stay home and watch TV, "as God meant us to." Yet now, with the ultraintelligent machines lying just below our horizon, it is time that we played this game in earnest, while we still have some control over the rules. In a few more years, it will be much too late.

Utopiamongering has been a popular and, on the whole, harmless occupation since the time of Plato; now it has become a matter of life and death-part of the politics of survival. Thinking machines, food production and population control must be considered as the three interlocking elements that will determine the shape of the future; they are not independent, for they all react on one another. This becomes obvious when we ask the question, which I have deliberately framed in as nonemotional a form as possible: "In an automated world run by machines, what is the optimum human population?"

There are many equations in which one of the possible answers is zero; mathematicians call this a trivial solution. If zero is the solution in this case, the matter is very far from trivial, at least from our self-centered viewpoint. But that it could-and probably will-be very low seems certain.

Fred Hoyle once remarked to me that it was pointless for the world to hold more people than one could get to know in a single lifetime. Even if one were President of United Earth, that would set the figure somewhere between 10.000 and 100,000; with a very generous allowance for duplication, wastage, special talents, and so forth, there really seems no requirement for what has been called the Global Village of the future to hold more than 1,000,000 people, scattered over the face of the planet. And if such a figure appears unrealistic-since we are already past the three-billion mark and heading for at least twice as many by the end of the century-it should be pointed out that once the universally agreed goal of population control is attained, any desired target can be reached in a remarkably short time. If we really tried (with a little help, perhaps, from the biology labs), we could reach a trillion within a century-four generations. It might be more difficult to go in the other direction, for fundamental psychological reasons, but it could be done. If the machines decide that more than 1,000,000 human beings constitutes an epidemic, they might order euthanasia for anyone with an I. Q. of less than 150, but I hope that such drastic measures will not be necessary.

Whether the population plateau levels off, a few centuries from today, at a million, a billion or a trillion human beings is of much less importance than the ways in which they will occupy their time. Since all the immemorial forms of "getting and spending" will have been rendered obsolete by the machines, it would appear that boredom will replace war and hunger as the greatest enemy of mankind.

One answer to this would be the uninhibited, hedonistic society of Huxley's Brave New World; there is nothing wrong with this, so long as it is not the only answer. (Huxley's unfortunate streak of asceticism prevented him from appreciating this point.) Certainly, much more time than at present will be devoted to sports, entertainment, the arts and everything embraced by the vague term "culture.

In some of these fields, the background presence of superior nonhuman mentalities would have a stultifying effect; but in others, the machines could act as pacemakers. Does anyone really imagine that when all the grand masters are electronic, no one will play chess? The humans will simply set up new categories and play better chess among themselves. All sports and games (unless they become ossified) have to undergo technological revolutions from time to time; recent examples are the introduction of fiberglass in pole vaulting, archery, boating. Personally, I can hardly wait for the advent of Marvin Minsky's promised robot table-tennis player.

These matters are not trivial: games are a necessary substitute for our hunting impulses, and if the ultraintelligent machines give us new and better outlets, that is all to the good. We shall need every one of them to occupy us in the centuries ahead.

Thinking machines will certainly make possible new forms of art and far more elaborate developments of the old ones. by introducing the dimensions of time and probability. Even today, a painting or a piece of sculpture that stands still is regarded as slightly passé. Although the trouble with most "kinetic art" is that it only lives up to the first half of its name, something is bound to emerge from present explorations on the frontier between order and chaos.

The insertion of an intelligent machine into the loop between a work of (continued on page 293)

MIND MACHINE

(continued from page 122) art and the person appreciating it opens up some fascinating possibilities. It would allow feedback in both directions: by this I mean that the viewer would react to the work of art, then the work would react to the viewer's reactions, then . . . and so on. for as many stages as was felt desirable. This sort of to-and-fro process is already hinted at, in a very crude way, with today's primitive "teaching machines"; and those modern novelists who deliberately scramble their text may also be groping in this direction. A dramatic work of the future, reproduced by an intelligent machine sensitive to the varying emotional states of the audience, would never have the same form, or even the same plot line, twice in succession. It would be full of surprises even to its human creatoror collaborator.

What sort of art machines would create for their own amusement and whether we would be able to appreciate it are questions that can hardly be answered today. The painters of the Lascaux Caves could not have imagined (though they would have enjoyed) the scores of art forms that have been invented in the 20,000 years since they created their masterpieces. Though in some respects we can do no better, we can do much more—more than any Paleolithic Picasso could possible have dreamed. And our machines may begin to build on the foundations we have laid.

Yet perhaps not. It has often been suggested that art is a compensation for the deficiencies of the real world: as our knowledge, power and, above all, our *maturity* increase, we will have less and less need for it. If this is true, the ultraintelligent machines would have no use for it at all.

Even if art turns out to be a dead end, there still remains science-the eternal quest for knowledge, which has brought man to the point where he may create his own successor. It is unfortunate that, to most people, "science" now means incomprehensible mathematical complexities; that it could be the most exciting and entertaining of all occupations is something that they find impossible to believe. Yet the fact remains that, before they are ruined by what is laughingly called education, all normal children have an absorbing interest and curiosity about the universe that, if properly developed, could keep them happy for as many centuries as they may wish to live.

Education: that, ultimately, is the key to survival in the coming world of thinking machines. The truly educated man (I have been lucky enough to meet two in my lifetime) can never be bored.



"... And these are my parents, Mr. and Mrs. Gerald N. LaVavooom."

The problem that has to be tackled within the next 50 years is to bring the entire human race, without exception, up to the level of semiliteracy of the average college graduate. This represents what may be called the *minimum* survival level; only if we reach it will we have a sporting chance of seeing the year 2200.

• • •

Perhaps we can now glimpse one viable future for the human race, when it is no longer the dominant species on this planet. As he was in the beg'nning, man will again be a fairly rare animal, and probably a nomadic one. There will be a few towns in places of unusual beauty or historical interest, but even these may be temporary or seasonal. Most homes will be completely self-contained and mobile, so that they can move to any spot on Earth within 24 hours.

The land areas of the planet will have largely reverted to wilderness: they will be much richer in life forms (and much more dangerous) than today. All adolescents will spend part of their youth in this vast biological reserve, so that they never suffer from that estrangement from nature that is one of the curses of our civilization.

And somewhere in the background in the depths of the sea, orbiting beyond the ionosphere—will be the culture of the ultraintelligent machines, going its own unfathomable way. The societies of man and machine will interact continuously but lightly: there will be no areas of conflict, and few emergencies, except geological ones (and these could be fully foreseeable). In one sense, for which we may be thankful, History will have come to an end.

All the knowledge possessed by the machines will be available to mankind, though much of it may not be understandable. There is no reason why this should give our descendants an inferiority complex; a few steps into the New York Public Library can do *that* 293 PLATBO

just as well, even today. Our prime goals will no longer be to discover but to understand and to enjoy.

Would the coexistence of man and machine be stable? I see no reason why it should not be, at least for many centuries. A remote analogy of this kind of dual culture-one society encapsulated in another-may be found among the Amish of Pennsylvania. Here is a selfcontained agricultural society, which has deliberately rejected much of the surrounding values and technology, yet is exceedingly prosperous and biologically successful. The Amish, and similar groups, are well worth careful study; they may show us how to get along with a more complex society that perhaps we cannot comprehend, even if we wish to.

For in the long run, our mechanical offspring will pass on to goals that will be wholly incomprehensible to us; it has been suggested that when this time comes, they will head on out into galactic space, looking for new frontiers, leaving us once more the masters (perhaps reluctant ones) of the Solar System, and not at all happy at having to run our own affairs. That is one possibility. Another has been summed up, once and for all, in the most famous short science-fiction story of our age. It was written by Fredric Brown almost 20 years ago, and it is high time that he received credit from the journalists who endlessly rediscover and quote him,

Fred Brown's story—as you have probably guessed—is the one about the supercomputer that is asked, "Is there a God?" After making quite sure that its power supply is no longer under human control, it replies in a voice of thunder: "Now there is."

This story is more than a brilliant myth; it is an echo from the future. For in the long run, it may turn out that the theologians have made a slight but understandable error—which, among other things, makes totally irrelevant the recent debates about the death of God.

It may be that our role on this planet is not to worship God—but to create him.

And then our work will be done. It will be time to play.

X



"One thing for sure . . . this isn't one of your run-of-the-mill Peeping Tom cases. . . ."

(continued from page 201) erotic art. Despite the persistence of book burning, even into our own time, the bolder and bawdier works of Aristophanes, Rabelais, Shakespeare, Ovid, D. H. Lawrence and others have survived countless suppressions; millions have read them, and every educated man at least knows that these giants of literature wrote such books. The average museumgoer, however, will be surprised to learn that similar works, in the visual rather than verbal medium, were produced by men such as Tintoretto, Rembrandt and Picasso.

That this revelation occasions shock stems from the fact that puritan taboos tend to linger longer over the "hot" visual arts than over the "cool" literary arts. Learning that Rembrandt portrayed sexual intercourse in realistic detail strikes many with the same impact as if they had been told that the Dutch master did the illustrations for a Mutt and Jeff eight-pager. It is to be hoped that this reaction is doomed and that it will vanish as the public begins to accept the fact that sexually explicit works, although under heavy ban, are a major part of man's artistic heritage. It is, quite obviously, as natural for a painter to explore sex frankly as it is for a poet, playwright or novelist.

Since we recently took a holiday from our own field of clinical psychology to arrange the first public showing in modern times of erotic painting and sculpture (held from May through July in Lund, Sweden; during September in Denmark: and scheduled for showing in Stockholm's Liljevalchs Museum from April 2, 1969, through May 1969), we are particularly pleased to write the text for this preview of what will be a PLAYBOY series of illustrated articles on artistic erotica, which will trace the history of this taboo-ridden art from the cave carvings of the Paleolithic Age up to the mixed-media experiments of today.

We ourselves see aesthetic, educational and even emotional benefits in visual erotica for the average art buff and museumgoer who has never been exposed to such material. He will acquire a new perspective on art; but even more importantly, he will realize that sexual subject matter can be presented with artistry, with beauty, sometimes with humor-and often with greatness. He will learn that sexual material need not be restricted to the sordid and amateurish treatment found in hard-core visual pornography; he will learn that sex has been the inspiration and sometimes the preoccupation of men of genius. This cannot help but allay anxieties, shame and misgivings that have been created by our society's traditional under-the-counter approach to crotica.