944 CLARKE, ARTHUR C. N. WIENER · MC 22

"The Evolutionary Cycle from Man to Machine" n.d.

The Evolutionary Cycle from Man to Machine

by ARTHUR C. CLARKE

ABOUT A MILLION YEARS AGO, an unprepossessing primate discovered that his forelimbs could be used for other purposes besides locomotion. Objects like sticks and sones could be grasped, and, once grasped, were useful for killing game, digging up roots, defending or attacking, and a hundred other jobs. On the third planet of the Sun, tools had appeared; and the place would never be the same again.

The first users of tools were *not* men—a fact appreciated only in the last year or two — but prehuman anthropoids; and by their discovery they doomed themselves. For even the most primitive of tools, such as a naturally-pointed stone that happens to fit the hand, provides a tremendous physical and mental stimulus to the user. He has to walk erect; he no longer needs huge canine teeth—since sharp flints can do a better job—and he must develop manual dexterity of a high order. These are the specifications of Homo sapiens; as soon as they start to be filled, all earlier models are headed for rapid obsolescence. To quote Professor Sherwood Washburn of the University of California's Anthropology Department: "It was the success of the simplest tools that started the whole trend of human evolution and led to the civilizations of today."

Note that phrase — "the whole trend of human evolution." The old idea that man invented tools is therefore a misleading half-truth; it would be more accurate to say that tools invented man. They were very primitive tools, in the hands of creatures who were little more than apes. Yet they led to us, and to the eventual extinction of the apemen who first wielded them.

Now the cycle is about to begin again; but neither history nor prehistory ever exactly repeats itself, and this time there will be a fascinating twist in the plot. The tools the apemen invented caused them to evolve into their successor, Homo sapiens. The tool we have invented *is* our successor. Biological evolution has given way to a far more rapid process—technological evolution. To put it bluntly and brutally, the machine is going to take over.

This, of course, is hardly an original idea. That the creations of man's brain might one day threaten and perhaps destroy him is such a tired old cliché that no self-respecting science-fiction magazine would care to use it. It goes back, through Capek's *R.U.R.*, Samuel Butler's *Erewhon*, Mary Shelley's *Frankenstein* and the Faust legend, to the mysterious but perhaps not wholly mythical figure of Daedalus, King Minos' one-man Office of Scientific Research. For at least three thousand years, therefore, a vocal minority of mankind has had grave doubts about the ultimate outcome of technology. From the self-centered, human point of view, those doubts are justified. But that, I submit, will not be the only—or even the most important—point of view for much longer.

When the first large-scale electronic computers appeared some fifteen years ago, they were promptly nicknamed "Giant Brains"—and the scientific community, as a whole, took a poor view of the designation. But the scientists objected to the wrong word. The electronic computers were not *giant* brains; they were dwarf brains, and they still are, though they have grown a hundred-fold within less than one generation of mankind. Yet even in their present flint-ax stage of evolution, they have done things which not long ago almost everyone would have claimed to be impossible — such as translating from one language to another, composing music, and playing a fair game of chess. And much more important than any of these infant *jeux* is the fact that they have breached the barrier between brain and machine.

This is one of the greatest—and perhaps one of the last — breakthought, like the discovery that the throughs in the history of human Earth moves round the Sun, or that man is part of the animal kingdom, or that $E = mc^2$. All these ideas took time to sink in, and were fanatically denied when first put forward. In the same way it will take a little while for men to realize that machines cannot only think, but may one day think men off the face of the Earth.

At this point you may reasonably ask: "Yes—but what do you mean by *think?*" I propose to side-step that question, using a neat device of the English mathematician A. M. Turing. Turing imagined a game played by two teletype operators in separate rooms — this impersonal link being used to remove all clues given by voice, appearance and so forth. Suppose that one operator was able to ask the other any questions he wished, and the other had to make suitable replies. If, after some hours or days of this conversation, the questioner could not decide whether his telegraphic acquaintance was human or purely mechanical, then he could hardly deny that he/it was capable of thought. An electronic brain that passed this test would, surely, have to be regarded as an intelligent entity. Anyone who argued otherwise would merely prove that he was less intelligent than the machine; he would be a splitter of nonexistent hairs, like the scholar who proved that the *Odyssery* was not written by Homer, but by another man of the same name.

We are still decades — but not centuries—from building such a machine, yet already we are sure that it could be done. If Turing's experiment is never carried out, it will merely be because the intelligent machines of the future will have better things to do with their time than conduct extended conversations with men. I often talk with my dog, but I don't keep it up for long.

The fact that the great computers of today are still high-speed morons, capable of doing nothing beyond the scope of the instructions carefully programed into them, has given many people a spurious sense of security. No machine, they argue, can possibly be more intelligent than its makers—the men who designed it, and planned its functions. It may be a million times faster in operation, but that is quite irrelevant. Anything and everything that an electronic brain can do must also be within the scope of a human brain, given sufficient time and patience. Above all, no machine can show originality or creative power or the other attributes which are fondly labeled "human."

The argument is wholly fallacious; those who still bring it forth are like the buggy-whip makers who used to poke fun at stranded Model Ts. Even if it were true, it could give no comfort, as a careful reading of these remarks by Dr. Norbert Wiener will show: "This attitude [the assumption that machines cannot possess any degree of originality] in my opinion should be rejected entirely. . . . It is my thesis that machines can and do transcend some of the limitations of their designers. . . . It may well be that in principle we cannot make any machine, the elements of whose behavior we cannot comprehend sooner or later. This does not mean in any way that we shall be able to comprehend them in substantially less time than the operation of the machine, nor even within any given number of years or generations. . . . This means that though they are theoretically subject to human criticism, such criticism may be ineffective until a time long after it is relevant."

In other words, even machines *less* intelligent than men might escape from our control by sheer speed of operation. And, in fact, there is every reason to suppose that machines will become much more intelligent than their builders, as well as incomparably faster.

There are still a few authorities who refuse to grant any degree of intelligence to machines, now or in the future. This attitude shows a striking parallel to that adopted by the chemists of the early Nineteenth Century. It was known then that all living organisms are formed from a few common elements—mostly carbon, hydrogen, oxygen and nitrogen —but it was firmly believed that the materials of life could not be made from "mere" chemicals alone. There must be some other ingredient — some essence or vital principle, forever unknowable to man. No chemist could ever take carbon, hydrogen and so forth and combine them to form any of the substances upon which life was based. There was an impassable barrier between the worlds of "inorganic" and "organic" chemistry.

This *mystique* was destroyed in 1828, when Wohler synthesized urea, and showed that there was no difference at all between the chemical reactions taking place in the body and those taking place inside a retort. It was a terrible shock to those pious souls who believed that the mechanics of life must always be beyond human understanding or imitation. Many people are equally shocked today by the suggestion that machines can think, but their dislike of the situation will not alter it in the least.

Since this is not a treatise on computer design, you will not expect me to explain how to build a thinking machine. In fact, it is doubtful if any human being will ever be able to do this in detail, but one can indicate the sequence of events that will lead from H. sapiens to M. sapiens. The first two or three steps on the road have already been taken; machines now exist that can learn by experience, profit from their mistakes and—unlike human beings — never repeat them. Machines have been built which do not sit passively waiting for instructions, but which explore the world around them in a manner which can only be called inquisitive. Others look for proofs of theorems in mathematics or logic, and sometimes come up with surprising solutions that had never occurred to their makers.

These faint glimmerings of original intelligence are confined at the moment to a few laboratory models; they are wholly lacking in the giant computers that can now be bought by anyone who happens to have a few hundred thousand dollars to spare. But machine intelligence will grow, and it will start to range beyond the bounds of human thought as soon as the second generation of computers appears—the generation that has been designed, not by men, but by other, "almost intelligent" computers. And not only designed, but also built—for they will have far too many components for manual assembly.

It is even possible that the first genuine thinking machines may be *grown* rather than constructed; already some crude but very stimulating experiments have been carried out along these lines. Several artifical organisms have been built which are capable of rewiring themselves to adapt to changing circumstances. Beyond this there is the possibility of computers which will start from relatively simple beginnings, be programed to aim at specific goals, and search for them by constructing their own circuits, perhaps by growing networks of threads in a conducting medium. Such a growth may be no more than a mechanical analogy of what happens to every one of us in the first nine months of our existence.

All speculations about intelligent machines are inevitably conditioned —indeed, inspired—by our knowledge of the human brain, the only thinking device currently on the market. No one, of course, pretends to understand the full workings of the brain, or expects that such knowledge will be available in any foreseeable future. (It is a nice philosophical point as to whether the brain can ever, even in principle, understand itself.) But we do know enough about its physical structure to draw many conclusions about the limitations of "brains" — whether organic or inorganic.

There are approximately ten billion separate switches—or neurons —inside your skull, "wired" together in circuits of unimaginable complexity. Ten billion is such a large number that, until recently, it could be used as an argument against the achievement of mechanical intelligence. About ten years ago a famous neurophysiologist made a statement (still produced like some protective incantation by the advocates of cerebral supremacy) to the effect that an electronic model of the human brain would have to be as large as the Empire State Building, and would need Niagara Falls to keep it cool when it was running.

This must now be classed with such interesting pronouncements as "No heavier-than-air machine will ever be able to fly." For the calculation was made in the days of the vacuum tube, and the transistor has now completely altered the picture. Indeed—such is the rate of technological progress today—the transistor itself is being replaced by still smaller and faster devices, based upon abstruse principles of quantum physics. If the problem were merely one of space, today's electronic techniques would allow us to pack a computer as complex as the human brain onto a single floor of the Empire State Building.

It's a tough job keeping up with science, and since I wrote that last paragraph the Marquardt Corporation's Astro Division has announced a new memory device which could store inside a six-foot cube *all information recorded during the last ten thousand years.* This means, of course, not only every book ever printed, but *everything* ever written in *any* language on paper, papyrus, parchment or stone. It represents a capacity untold millions of times greater than that of a single human memory, and though there is a mightly gulf between merely storing information and thinking creatively —the Library of Congress has never written a book — it does indicate that mechanical brains of enormous power could be quite small in physical size. And the shrinkage is just gaining momentum, if I may employ such a mind-boggling phrase. Westinghouse now manufacturers a five-watt amplifier that could rather easily be mistaken for an aspirin tablet, and radio sets the size of lumps of sugar are also available. Before long, they will be the size not of lumps but of grains, for the slogan of the microminiaturization experts is "If you can see it, it's too big."

Just to prove that I am not exaggerating, here are some statistics you can use on the next hi-fi fanatic who takes you on a tour of his installation. During the 1950s, the electronic engineers learned to pack up to a hundred thousand electronic components into one cubic foot. (To give a basis of comparison, a good hi-fi amplifier may contain two or three hundred parts, a domestic radio about a hundred.) Here at the beginning of the Sixties, the attainable figure is around a million electronic components per cubic foot; by 1970, when today's experimental techniques of microscopic engineering have begun to pay off, it may reach a hundred million.

Fantastic though this last figure is, the human brain surpasses it by a thousand-fold, packing its ten billion neurons into a *tenth* of a cubic foot. And although smallness is not necessarily a virtue, even this may be nowhere near the limit of possible compactness.

For the cells composing our brains are slow-acting, bulky and wasteful of energy — compared with the scarecly more than atom-sized computer elements that are theoretically possible. The mathematician Jhon von Neumann once calculated that electronic cells could be ten billion times more efficient than protoplasmic ones; already they are a million times swifter in operation, and speed con often be traded for size. If we take these ideas to their ultimate conclusion, it appears that a computer equivalent in power to one human brain need be no bigger than a matchbox.

This slightly shattering thought becomes more reasonable when we take a critical look at flesh and blood and bone as engineering materials. All living creatures are marvelous, but let us keep our sense of proportion. Perhaps the most wonderful think about life is that it works at all, when it has to employ such evtraordinary materials, and has to tackle its problems in such roundabout ways.

As a perfect example of this, consider the eye. Suppose *you* were given the problem of designing a camera—for that, of course, is what the eye is—which *has to be constructed entirely of water and jelly,* without using a scrap of glass, metal or plastic. You're quite right; the feat is impossible. The eye is an evolutionary miracle, but it's a lousy camera. You can prove this while you're reading the nevt sentence. Here's a medium-length word: photograph. Close one eye and keep the other fixed—repeat, *fixed*—on that center "g." You may be surprised to discover that, unless you cheat by altering the direction of your gaze, you cannot see the whole word clearly. It fades out three or four letters to the right and left.

No camera built—even the cheapest—gives as poor an optical performance as this. For color vision, also, the human eye is nothing to boast about; it can operate only over a small band of the spectrum. To the worlds of the infrared and ultraviolet, visible to bees and other insects, it is completely blind.

We are not conscious of these limitations because we have grown up with them, and indeed if they were corrected the brain would be quite unable to handle the vastly increased flood of information. But let us not make a virtue of a necessity; if our eyes had the optical performance of even the cheapest miniature camera, we would live in an unimaginably richer and more colorful world.

These defects are due to the fact that precision scientific instruments simply cannot be manufactured from living materials. With the eye, the ear, the nose—indeed, all the sense organs — evolution has performed a truly incredible job against fantastic odds. But it will not be good enough for the future; indeed, it is not good enough for the present.

There are some senses that do not exist, that can probably never be provided by living structures, and which we need in a hurry. On this planet, to the best of our knowledge, no creature has ever developed organs that can detect radio waves or radioactivity. Though I would hate to lay down the law and claim that nowhere in the universe can there be organic Geiger counters or living TV sets, I think it highly improbable. There are some jobs that can be done only by vacuum tubes or magnetic fields or electron beams, and are therefore beyond the capability of purely organic structures.

There is another fundamental reason why living machines such as you and I cannot hope to compete with nonliving ones. Quite apart from our poor materials, we are handicapped by one of the toughest engineering specifications ever issued. What sort of performance would you expect from a machine which has to grow several billionfold during the course of manufacture, and which has to be completely and continuously rebuilt, molecule by molecule, every few weeks? This is what happens to all of us, all the time; you are not the man you were last year, in the most literal sense of the expression.

Most of the energy and effort required to run the body goes into its perpetual tearing down and rebuilding, a cycle completed every few weeks. New York City, which is a very much simpler structure than a man, takes hundreds of times longer to remake itself. When one tries to picture the body's myriads of Conrad Hiltons and Bill Zeckendorfs all furiously at work, tearing up arteries and nerves and even bones, it is astonishing that there is any energy left over for the business of thinking.

Now I am perfectly well aware that many of the "limitations" and "defects" just mentioned are nothing of the sort, looked at from another point of view. Living creatures, because of their very nature, can evolve from simple to complex organisms. They may well be the only path by which intelligence can be attained, for it is a little difficult to see how a lifeless planet can progress directly from metal ores and mineral deposits to electronic computers by its own unaided efforts. But though intelligence can only arise from life, it may then discard it. Perhaps at a later stage, as the mystics have suggested, it may also discard matter; but this leads us into realms of speculations which an unimaginative person like myself would prefer to avoid.

One often-stressed advantage of living creatures is that they are self-repairing and can reproduce themselves with ease, indeed, with enthusiasm. This superiority over machines will be short-lived; the general principles underlying the construction of self-repairing and self-reproducing machines have already been worked out. There is, incidentally, something ironically appropriate in the fact that A. M. Turing, who pioneered in this field and first indicated how thinking machines might be built, shot himself a few years after publishing his results. It is very hard not to draw a moral from this.

The greatest single stimulus to the evolution of mechanical — as opposed to organic — intelligence is the challenge of space. Only a vanishingly small fraction of the universe is directly accessible to mankind, in the sense that we can live there without elaborate protection or mechanical aids. If we generously assume that humanity's potential *Lebensraum* extends from sea level to a height of three miles, over the whole Earth, that gives us a total of some half billion cubic miles. At first sight this is an imressive figure, but it is absolutely nothing when set against the reaches of space. Our present telescopes, which are certainly not the last word on the subject, sweep a volume at least a million million million million million million million million million million times greater. Though such a number is, of course utterly beyond conception, it can be given a vivid meaning. If we reduce the known universe to the size of the Earth, then the portion in which *we* can live without space suits and pressure cabins is about the size of a single atom.

It is true that, one day, we are going to explore and colonize many other atoms in this Earth-sized volume, but it will be at the cost of tremendous technical efforts, for most of our energies will be devoted to protecting our frail and sensitive bodies against the extremes of temperature, pressure or gravity found in space and on other worlds. Within very wide limits, machines are indifferent to these extremes. Even more important, they can wait patiently through the years and the centuries that will be needed for travel to the far reaches of the universe.

Creatures of flesh and blood such as ourselves can explore space and win control over infinitesimal fractions of it. But only creatures of metal and plastic can ever really conquer it, as indeed they have already started to do. The tiny brains of our Vanguards and Explorers and Pioneers barely hint at the mechanical intelligences that will one day be launched at the stars.

It may well be that only in space, confronted with environments fiercer and more complex than any to be found upon this planet, will intelligence be able to reach its fullest stature. Like other qualities, intelligence is developed by struggle and conflict; in the ages to come, the dullards may remain on placid Earth, and real genius will flourish only in space—the realm of the machine, not of flesh and blood.

A striking parallel to this situation can already be found on our planet. Some millions of years ago, the most intelligent of the mammals withdrew from the battle of the dry land and returned to their ancestral home, the sea. They are still there, with brains larger and potentially more powerful than ours. But (as far as we know) they do not use them; the static environment of the sea makes little call upon intelligence. The porpoises and whales, who might have been our equals and perhaps our superiors had they remained on land, now race in simple-minded and innocent ecstasy beside nuclear-powered sea monsters carrying sixteen megatons of death. Perhaps they, not we, made the right choice, but it is too late to join them now.

If you have stayed with me so far, the protoplasmic computer inside your skull should now be programed to accept the idea—at least for the sake of argument—that machines can be both more intelligent and more versatile than men, and may well be so in the very near future (probably before the end of the next century; no one can imagine *any* technical development that will take much longer than that). So it is time to face the question: "Where does that leave man?"

I suspect that this is not a question of very great importance—except, of course, to man. Perhaps the Neanderthalers made similar plaintive noises, around 100,000 B.C., when H. sapiens appeared on the secne, with his ugly vertical forehead and ridiculous protruding chin. Any Paleolithic philosopher who gave his colleagues the right answer would probably have ended up in the cooking-pot; I am prepared to take that risk.

The short-term answer may indeed be cheerful rather than depressing. There may be a brief Golden Age when men will glory in the power and range of their new partners. Barring war, this age lies directly ahead of us. As Dr. Simon Ramo put it recently: "The extension of the human intellect by electronics will become our greatest occupation within a decade." That is undoubtedly true, if we bear in mind that at a somewhat later date the word "extension" may be replaced by "extinction."

One of the ways in which thinking machines will be able to help us is by taking over the humbler tasks of life, leaving the human brain free to of course, that this is any guarantee concentrate on higher things. (Not, that it will do so.) For a few generations, perhaps, every man will go through life with an electronic companion, which may be no bigger than today's transistor radio. It will "grow up" with him from infancy, learning his habits, his business affairs, taking over all the dull chores like routine correspondence and income tax returns. On occasion it could even take its mater's place, keeping appointments he preferred to miss, and then reporting back in as much detail as he desired. It could substitute for him over the telephone so completely that no one would be able to tell whether man or machine was speaking: a century from now, Turing's "game" may be an integral part of our social lives, with complications and possibilities which I leave to the imagination.

You may remember that delightful robot, Robbie, from the movie *Forbidden Planet* (one of the three or four movies so far made that anyone interested in science-fiction can point to without blushing). I submit, in all seriousness, that most of Robbie's abilities—together with those of a better-known character, Jeeves—will one day be incorporated in a kind of electronic companion-secretary-valet. It will be much smaller and neater than the walking jukeboxes which Hollywood presents, with typical lack of imagination, when it wants to portray a robot. And it will be extremely talented, with quick-release connectors allowing it to be coupled to an unlimited variety of sense organs and limbs. It would, in fact, be a kind of general-purpose, disembodied intelligence that could attach itself to whatever tools were needed for any particular occasion. On one day it might be using microphones or electric typewriters or TV cameras; on another, automobiles or airplanes—or the bodies of men and animals.

And this is, perhaps, the moment to deal with a conception which many people find even more horrifying than the idea that machines will replace or supersede us. It is the idea that they may combine with us.

I do not know who first thought of this; probably the physicist J. D. Bernal, who in 1929 published an extraordinary book of scientific predictions called *The World, the Flesh and the Devil*. In this slim and long-out-of-print volume (I sometimes wonder what the sixty-year-old Fellow of the Royal Society now thinks of his youthful indiscretion, if he ever remembers it), Bernal decided that the numerous limitations of the human body could be overcome only by the use of mechanical attachments or substitutes —until, eventually, all that might be left of man's original organic body would be the brain.

This idea is already far more plausible than when Bernal advanced it, for in the last few decades we have seen the development of mechanical hearts, kidneys, lungs and other organs, the wiring of electronic devices directly into the human nervous system.

Olaf Stapledon developed this theme in his wonderful history of the future, *Last and First Men*, imagining an age of immortal Giant Brains, many yeards across, living in beehive-shaped cells, sustained by pumps and chemical plants. Though completely immobile, their sense organs could be wherever they wished, so their center of awareness — or consciousness, if you like — could be anywhere on Earth or in the space above it. This is an important point which we—who carry our brains around in the same fragile structure as our eyes, ears and other sense organs, often with disastrous results—may easily fail to appreciate. Given perfected telecommunications, a fixed brain is no handicap; but rather the reverse. Your present brain, totally imprisoned behind its walls of bone, communicates with the outer world and receives its impressions of it over the telephone wires of the central nervous system—wires varying in length from a fraction of an inch to several feet. You would never know the difference if those "wires" were actually hundreds or thousands of miles long, or included mobile radio links, and your brain never moved at all.

In a crude way—yet one that may accurately foreshadow the future—we have already extended our visual and tactile senses away from our bodies. The men who now work with radioisotopes, handling them with remotely-controlled mechanical fingers and observing them by television, have achieved a partial separation between brain and sense organs. They are in one place, their minds effectively in another.

Recently the word Cyborg (cybernetic organism) has been coined to describe a machine-animal of the type we have been discussing. Doctors Manfred Clynes and Nathan Kline of Rockland State Hospital, Orangeburg, New York, who invented the name, define a Cyborg in these stirring words: "an exogenously extended organizational complex functioning as a homeostatic system." To translate, this means a body which has machines hitched to it, or built into it, to take over or modify some of its functions.

I suppose one could call a man in an iron lung a Cyborg, but the concept has far wider implications than this. One day we may be able to enter into temporary unions with any sufficiently sophisticated machines, thus being able not merely to control but to *become* a spaceship or a submarine or a TV network. This would give far more than purely intellectual satisfaction; the thrill that can be obtained from driving a racing car of flying an airplane may be only a pale ghost of the excitement our great-grandchildren may know, when the individual human consciousness is free to roam at will from machine to machine, through all the reaches of sea and sky and space.

But how long will this partnership last? Can the synthesis of man and machine ever be stable, or will the purely organic component become such a hindrance that it has to be discarded? If this eventually happens—and I have tried to give reasons why it must — we have nothing to regret, and certainly nothing to fear.

The popular idea, fostered by comic strips and the cheaper forms of science-fiction, that intelligent machines must be malevolent entities hostile to man, is so absurd that it is hardly worth wasting energy to refute it. I am almost tempted to argue that only *un*intelligent manchines can be malevolent; anyone who has tried to start a balky outboard will probably agree. Those who picture machines as active enemies are merely projecting their own aggressive instincts, inherited from the jungle, into a world where such things do not exist. The higher the intelligence, the greater the degree of cooperativeness. If there is ever a war between men and machines, it is easy to guess who

Yet however friendly and helpful

will start it.

the machines of the future may be, most people will fee that it is a rather bleak prospect for humanity if it ends up as a pampered specimen in some biological museum—even if that museum is the whole planet Earth. This, however, is an attitude I find impossible to share.

No individual exists forever; why should we expect our species to be immortal? Man, said Nietzsche, is a rope stretched between the animal and the superhuman—a rope across the abyss. That will be a noble purpose to have served. ■