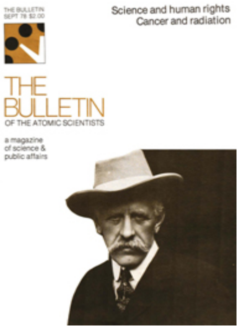

Bulletin of the Atomic Scientists

ISSN: 0096-3402 (Print) 1938-3282 (Online) Journal homepage: http://www.tandfonline.com/loi/rbul20

Once more—a computer revolution

Joseph Weizenbaum

To cite this article: Joseph Weizenbaum (1978) Once more—a computer revolution, Bulletin of the Atomic Scientists, 34:7, 12-19, DOI: 10.1080/00963402.1978.11458531

To link to this article: https://doi.org/10.1080/00963402.1978.11458531

Published online: 15 Sep 2015.

Much depends on knowing what limits to impose on the application of computers to human affairs and on knowing the impact of the computer on human dignity.

JOSEPH WEIZENBAUM

Once more—a computer revolution

The cost and the physical size of computer hardware are decreasing exponentially. Therefore, the very measures classifying computers as micro, mini, small, large and very large are also constantly being readjusted. Today's so-called minis are functionally equivalent, at least roughly, to the last decade's most powerful large computers, yet almost all their physical indices are a very small fraction of the corresponding indices of their ancestors. So is their cost. Moreover, computers, as large as the older computers, have many times as many components packed into them and are therefore functionally much more powerful while being no more expensive. This phenomenon is, of course, also reflected on the software side: the measures by which programs are ranged from small to large are changing similarly. Programs which only a few years ago would have been classified as rather large are now shoe-horned into mini and even micro computers while, on the other end of the scale, programs of hitherto unimagined size *and complexity* are developed for the newer giant computer systems.

Technological optimists believe that these dramatic developments *must inevitably* flood the marketplace with computers. Then, in the not too distant future, computers will pervade the American scene as, for example, fractional horsepower electric motors do today. Moreover like the small engines, these computers will be used in homes as unobtrusive parts of a wide variety of household gadgets.

The widely shared belief in technological inevitability, especially as it applies to computers, is translated by scholars and the popular media alike into the announcement of still another *Computer Revolution.* This much heralded revolution will supposedly transform society to its very core; the new Information Society will emerge.

The question asked only rarely is what pressing problems this inundation of technological fixes should attack.

Certain problem areas *are* often identified, to be sure. The drabness of modern society, for example, is sometimes discussed. Some computer scientists believe this drabness is largely due to the deadly uniformity of most consumer goods [1]. Versatile manufacturing robots, they argue, could relieve this monotony by "individualizing" products. Education is also raised as a problem area. Here it is occasionally argued [2] that the visions of such thinkers as Dewey, Montessori, and Neill "fail in practice *for lack of a technological basis.* The computer now provides it" [emphasis mine].

But in these and other cases the discussion is clearly carried out in a mode of thought that has become altogether too traditional: it begins with a great many solutions and then looks for problems. This way of thinking obscures real problems. The aimlessness of everyday life for millions in modern society is rooted in the individual's alienation from nature, work and other human beings. To give everyone (who can afford it) a pair of shoes superficially different from everyone else's and then advertise this as a step toward the amelioration of society is not revolutionary; it is absurd. Similarly, the real problems which people like Neill and Montessori actually confront are not functions of some "technological base." No fix of the American education system that does not recognize that American schools are rapidly becoming America's principal juvenile minimum security prisons can have socially therapeutic effects.

Home computer enthusiasts struggle with problems that could arise only from the triumph of mass marketing techniques such as those which gave us the multi-million dollar feminine-hygiene (that is, vaginal) deodorant industry. The product to be marketed is invented simultaneously with the dysfunction designed to be cured. It is assumed that, as prices are lowered below any conceivable threshold of consumer resistance, virtually every household will have a programmable computer. The *problem* then is created by the *solution:* what are people to do with this appliance; what should it be applied *to?*

A typical essay on the home computer begins by assuming that there *are* computers in the home and then questions what they may be used for. The home computers foreseen are (in miniature, to be sure) like those existing in the world at large—including, for example, freestanding computers on which anything at all may be programmed, computers equipped with prepackaged systems, and process control computers. The issues that emerge from considerations of the home computer are much the same as those which arise from the presence of the computer in modern society generally. Chief among these are questions which probe what social needs computers help satisfy today and what roles they are likely to play tomorrow.

Perhaps the first question is just what fraction of American homes will have the kinds of computing machines typically envisioned. A standard analogy is to television. Essentially *all* American dwellings have at least one television set. In-

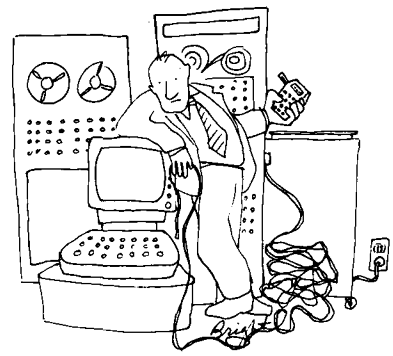

deed, many dwellings of the poor and the very poor have, whatever else they lack, a television set. Television is an example of a technological gadget which vindicated the marketeers who think in terms of consumer resistance thresholds below which it is possible to duck absolutely. Nevertheless, the poor often purchase television sets at a cost that is outrageous when measured in terms of elementary necessities given up.

Will the home computer be as pervasive as today's television sets? The answer almost certainly is *no*. The home, pictured in the accounts of home computer advocates [3], is middle class, or even upper middle class. There are some appliances computers must control: the wall-to-wall carpeting must be cleaned by a robot, roasts are in the oven, the computer helps "the mother" pay the telephone bill, and so on and on. Another computer scientist, B. O. Evans [4], imagines the same kind of home when he addresses himself to the home computer. He writes:

> For home use terminals have potential for catalog ordering, activity planning, home library and education, family health including histories, diagnoses, doctors' specialty lists, emergency procedures; family recreation including music selection and games; career guidance, tax records and returns, home safety and property maintenance including house plan retrieval, maintenance schedules, electrical and other physical facility layouts and energy management; budgeting and banking.

What and whose *needs* will these functions and the ongoing proliferation of computers and computer controlled systems satisfy? What will be the indirect effects on a society that is monitored and controlled by systems even its own technostructure little understands?

We may recall the euphoric dreams that were articulated by then Secretary of Commerce Herbert Hoover, at the dawn of commercial radio broadcasting, and again by others when mass television broadcasting was about to become a reality. It was foreseen that these media would exert an enormously beneficial influence on American culture. Americans would be exposed to correct English, to great literature, great drama, and so on.

We all know what actually happened. The technological dream was more than realized, but the cultural dream was not. Our magnificent technology, more than Wagnerian in its proportions, that combines the technology of precise guidance of rockets, of space flight, of the most clever and intricate electronics, of photography, that exquisitely refines a combination of some of the human species' highest intellectual achievements—delivers what to the masses? An occasional gem buried in immense floods of the most banal and insipid or pathological elements in our civilization.

In its current form, the home computer is merely a miniature version of the free-standing computers which can be found in countless laboratories, business offices and other enterprises. However, just as many of these computers are being increasingly interconnected to one another to form computer networks so, according to most authorities, will home computers become satellites of a variety of large computer networks. Only in this way would the home computers be able to access the large data bases required for what Evans visualized as these computers' function. Indeed, many authorities believe that home computers which function as nodes of extensive computer networks will play a crucial part in the process of transforming our society into what Daniel Bell, professor of sociology at Harvard University, calls an *Information Society.*

Bell sees the Information Society as a child of the marriage of modern communication and computer technologies. Certainly, one foundation of the Information Society is knowledge, which Bell defines as

> . . . an organized set of statements of facts or ideas, presenting a reasoned judgment or an experimental result which is transmitted to others through some communication medium in some systematic form. Thus I distinguish knowledge from news or entertainment. Knowledge consists of new judgments (research and scholarship) or presentations of older judgments (textbook, teaching and library and archive materials) [5].

He characterizes this definition as an attempt at an " 'objective definition' that would allow a researcher to plot the growth and use of knowledge."

What renders Bell's definition of knowledge nearly useless for the present purpose is that it is fatally circular and incomplete: What *facts, experimental results,* and *reasoned judgments* are is itself determined by the observer's organizing principles. This, the observer's *Weltanschauung,* however, is itself largely tacit. Bell's definition is incomplete also because it systematically excludes almost everything called knowledge in everyday life. People know a great many things that are neither products of research and scholarship nor materials in textbooks or archives. They know what pleases people they see every day and what offends

Joseph Weizenbaum is professor of computer science at the Massachusetts Institute of Technology. He is the author of *Computer Power and Human Reason* (1976).

them. They know their way about their cities and what detours to take when the usual paths are blocked.

Bell's willingness to exclude this kind of knowledge from consideration betrays the kind of parochialism that afflicts almost the entire intelligentsia. It betrays what the intelligentsia counts, not only as knowledge, but as fact. Bell himself hints what these determinants are—and what they are not:

> The upheaval in telecommunications and knowledge poses two policy problems, one 'structural,' the other, intellectual. . . .
>
> The structural question is what kind of technical-economic organization is best designed to be efficient, meet consumer *(that is, industrial, commercial, financial, scientific, library)* use, and *allow for continuing technological development* without becoming overly rigidified. . . .
>
> The second, intellectual-political, problem is the question of a national information policy, *particularly the dissemination of science and technical information* (emphasis mine) [5].

For Bell, "the crucial variables of the post-industrial society are information and knowledge." And Bell with obvious approval attributes to the psychologist George A. Miller the observation that:

> . . . recoding is an extremely powerful weapon for increasing the amount of information that we can deal with. In one form or another we use recoding constantly in our daily behavior. . . .
>
> Our language is tremendously useful for repackaging material into a few chunks rich in information . . . the kind of linguistic recoding people do seems to me to be the very lifeblood of the thought process [5].

I agree entirely.

What Miller speaks of as chunking is the phenomenon which permits us to recall, say, the telephone area code of New York, not as the sequence of the three separate integers 2 and 1 and 2 but as the single number 212. Even more importantly, words like *mother, enemy,* and so on, are not remembered merely as words; but as chunks which engage huge, often conflicting, conceptual structures laden with emotional significance. A welfare computer system may very well be able to dip into its data bank and calculate that, say, five people occupy a particular household; it cannot understand what difference it makes whether those five people are merely roomers who happen to share the rent-burden, or a family. Still, the human meaning of a family cannot be part of the computed chunk.

Computer-based information systems *necessarily* induce recoding of data into information-rich chunks. But the recoding required for the computer denudes the original data of the subtleties which accompanied them and determined their meanings while still in ordinary language. The *richness* of chunks created either by or for the computer is of a different order than that of their sources. As Bell himself says:

> . . . if the purpose of a library, or of a knowledge-base computer program, is to help an historian to assemble evidence, or a scholar to "reorder" ideas, then the very ambiguity of language, reflecting the fact that terms necessarily vary in different contexts and lend themselves to different interpretations, or the shifting historical usages over time ... makes the problem of a "computer knowledge" program quite different from an "information program". . . . A sophisticated reader, studying a philosophical text, may make use of the existing index at the back of the book, but if he is to absorb, and use the ideas in a different and creative way, he has to, necessarily, create his own index by re-grouping and re-categorizing the terms he has employed. . . . *And in this process, no mechanical ordering, no exhaustive set of permutations and combinations can do the task* (emphasis mine) [5].

This passage voices Bell's conviction that there are limits to what computers can do, particularly that artificial intelligence cannot produce an artifact that exhibits the entire range of human creativity. In this, we agree. Bell, however, seems to see this boundary as being relevant to only the most extreme fantasies of the leadership of the artificial intelligence community (the artificial intelligentsia), hence irrelevant to practical current concerns. This is where he and I disagree.

The use of large-scale computer based information systems induces an extremely poverty stricken notion of knowledge and fact. Unfortunately, this same notion—a kind of pragmatic positivism bordering on scientism—dominates much of the thinking of modern intellectuals and political leaders, as well as ordinary people. It has no *necessary* relationship to the computer, but the computer is its most stark symbolic manifestation.

To see its influence one may turn to Bell's own examples. Consider first the report of the Club of Rome, the *Limits to Growth* study, about which Bell writes:

> What gave the Club of Rome study a degree of authority was the announcement that the authors had succeeded in "modeling" the world economy and carrying out a computer simulation which traced out the interconnection of four basic variables: resources, population, industrial production and pollution [5].

Bell goes on to remark that "the *Limits to Growth* study has been largely discredited." Nevertheless, the study had and continues to have *authority.* It was not the announcement, however, that lent authority to the study; it was the fact that the study was conducted by insiders at the Massachusetts Institute of Technology, *and* that the model being

Decisions crucially affecting people's lives are made with the aid of computer systems contaminated by a "broad spectrum of inadequate and inaccurate economic indices" and by systematic lies.

announced was done on a computer.*

Interestingly enough, Bell claims the *Limits to Growth* model was discredited by the "unreliability of [its] initial data" and by "its simplified assumption of a linear, extrapolative growth." He never hints that this or any other model's difficulties might be in their epistemological foundations.

Jay Forrester has repeatedly revealed his models' epistemological foundations. For example:

> . . . the human mind is not adapted to interpreting how social systems behave . . . until recently there has been no way to estimate the behavior of social systems except by contemplation, discussion, argument, and guess work.
> . . . The great uncertainty with mental models is the inability to anticipate consequences of interactions between parts of a system. This uncertainty is *totally* eliminated in computer models. Given a stated set of assumptions, the computer traces the resulting consequences without doubt or error. . . . Furthermore, any concept or relationship that can be clearly stated in ordinary language can be translated into computer model language [6].

The widely shared belief in the *epistemology* expressed by these words is chiefly responsible for the acceptance of Forrester's and similar models.

Consider the impact of Forrester's words. Those he addresses hear that their thinking leads to uncertainty whereas Forrester-like computer models totally eliminate uncertainty, doubt or error. That is what they *hear;* it is not *precisely* what Forrester said. For he said only that given a system of well-formed equations, their solutions (if they exist) are unambiguously determined. With that one cannot quarrel. But the word *doubt* is curiously out of place. It is a word out of psychology, not mathematics or logic. What Forrester means is that because of inherent uncertainty, one must doubt conclusions reached from mere thinking. Conclusions derived from computer models are valid beyond doubt.

Forrester's stated *assumptions* may be correct or they may be incorrect, but they must necessarily be incomplete. And their necessary incompleteness derives from exactly the same source as the incompleteness of Bell's set of knowledge. The last sentence of the above quotation implies that anything worth saying at all—hence worth knowing in Bell's sense—can be "clearly stated in ordinary language," and hence "translated to computer model language."* Bell rests his vision of the coming Information Society on precisely this epistemological foundation. For "the crucial variables of the post-industrial society," Bell argues,

> . . . are information and knowledge. By information I mean, in the broadest sense, data processing. And the storing, retrieval and processing of data becomes the essential resource for all economic and social exchanges [5].

He envisions an extension of what already exists, namely the widespread use of computers to do the data processing in the information and knowledge society. Furthermore, almost all the processing will be done on data bases also stored in computer systems. These Bell characterizes as "characteristics of populations: census data, market research, opinion surveys, election data, etc." But what about these data bases? Bell himself quotes Peter H. Schuck:

> What is . . . disturbing, given the imminence of national economic planning, is the abject poverty of our economic statistical base, upon which a good theory must be grounded. In recent years the inadequacy and inaccuracy of a broad spectrum of economic indices—including the wholesale price index, the consumer price index, the unemployment rate, and business inventory levels—have become quite evident [5].

While the computer induces confidence (as in the *Limits to Growth* study), it usually magnifies errors and their consequences enormously.

Another classic example comes from the much-touted command and control system in operation during the Vietnam war. "The mechanisms of [this system] were so complete," Bell says, "that basic tactical decisions . . . were controlled by political centers in the White House, ten thousand miles away, but transmitted in "real time." However *complete* this system may have been, Admiral Moorer (then Chairman of the Joint Chiefs of Staff) testified to the U.S. Senate Armed Services Committee that specially programmed computers in the field systematically lied to the Pentagon's computers about the secret bombing of Cambodia.

It is instructive to note just how the U.S. Air Force computers in Vietnam were made to lie to the Pentagon's computers in Washington. Computers in the field were programmed automatically to convert the geographical coordinates of targets struck by U.S. planes in Cambodia to coordinates of *legitimate* targets in Vietnam. Tapes of these allegedly *raw,* though actually *cleansed,* data were then forwarded to Washington to be entered into the Pentagon's computers.

From a military point of view, this raises serious questions of command and control. It raises even more general questions of responsibility and accountability. The relatively simple technical task of writing the coordinate conversion programs had to be assigned by someone to someone. Perhaps the programmers who actually did the job were given their assignment in purely abstract form,

*M.I.T. proudly characterizes itself as being "polarized around science and technology."

*In these circumstances one needs to recall Eugene Ionesco's remark: "not everything is unsayable in words—only the living truth."

without being told what the ultimate function of their product was. If so, then the programmers could deny responsibility for the consequences of their handiwork on the grounds that they didn't know what they were doing. On the other hand, perhaps they knew what they were doing but, being in the military, thought it their duty to follow orders and, more importantly, perhaps they felt that that duty removed all responsibility.

Decisions crucially affecting people's lives are made with the aid of computer systems contaminated by a "broad spectrum of inadequate and inaccurate economic indices" and by systematic lies. If the programmers of these systems—and by extension, their professional managers, systems analysts, etc.—are not responsible for the consequences of actions based on what these computer systems tell policy-makers, and if policy-makers are excused from responsibility on the grounds that they merely relied on "what the computer said," then who is responsible?

This question poses a special case of the problem of individual responsibility and accountability that has manifested itself most egregiously in Adolf Eichmann's claim of personal innocence. Is there any *moral* difference between Eichmann's failure to confront what he was actually doing and the Air Force programmers' identical failure?

We Americans puzzle over the circumstance that neither General Westmoreland nor Lieutenant Calley are responsible for the men and women and babies Calley shot and killed with his own hand. Nor do any individuals in the U.S. military's chain of command acknowledge accountability for implementing the disinformation machinery—to borrow a truly Orwellian term from the world's intelligence agencies—which systematically deceived at least some policy-makers. There appear to be no actors on stage, only anonymous events.

The political and social institutionalization of systematic retreats from responsibility and accountability has no *necessary* relation to computers. However, the computer, and particularly the role advocated for it by many social scientists and computer intellectuals, amplifies and intensifies the problem and exacerbates its effects.

Computer intellectuals are aware of this and sometimes voice their concern—but usually in ways oddly detached from present-day reality.

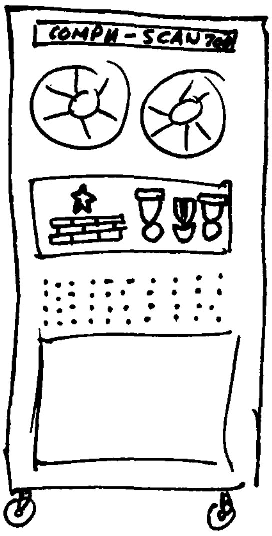

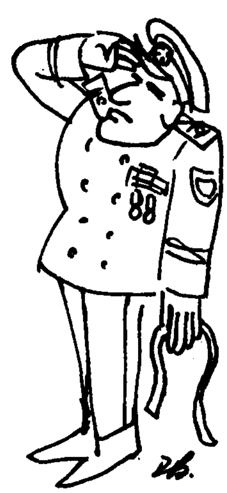

For example, Alan Perlis, head of the Computer-Science Department at Yale University, sees the computer as "having a day to day effect on man and his society," as "pulsing an 'ecological' transformation." Perlis "sees [computer science] as studying the nature and consequences of the phenomena arising around, and because of, computers" [7]. Yet, by far the most important of these phenomena—the "transformations in man and society" induced by the computer—is strangely absent from his agenda of frontier research.

However, Perlis isn't totally unaware of some social and political problems sharpened by the application of certain computer technologies. For example, he acknowledges that "research on speech understanding [by computers] can lead to programs that eavesdrop or deny us human contact in some telephone regulated transactions." But then he dismisses the crucial problems of responsibility and control, by saying: "The programs are not required to either exploit or correct our social deficiencies." I suppose this is intended to absolve *programs* from any responsibility for any harm that may come from their use—just as bullets are not responsible for the people they kill. But not a word about the responsibilities of the researchers.

To give another example, Marvin Minsky, Donner Professor of Science at M.I.T., confesses that he is

> . . . inclined to fear most the HAL scenario [referring to the computer on board the spaceship in the movie *2001* which eventually wrested control from the ship's astronauts]. The first AI [artificial intelligence] systems of large capability will have many layers of poorly understood control-structure, and obscurely encoded goal-structure: If it cannot edit its high-level intentions, it may not be smart enough to be useful, but if it can, how can the designers anticipate the machine it evolves into? In a word, I would expect the first self-improving AI machines to become "psychotic" in

Are there technical solutions to the problem presented by essentially incomprehensible computer systems? I don't think so.

many ways, and it may take many generations to "stabilize" them. The problem could become serious if economic incentives to use early unreliable systems are large—unfortunately there are too many ways a dumb system with a huge data base can be useful [8].

Minsky believes himself to be talking about machines of the future. His *fear* is therefore abstract and has little, if any, influence on what he believes he ought to worry about today. But if we, as we should, conceive of computer systems as including the people who manage and maintain them, then it becomes clear that the "early unreliable systems" Minsky rightly fears are already very much with us and that the economic incentives to use them are, for many organizations, already insuperably large.

Almost all very large computer systems *have* "many layers of poorly understood control-structure and poorly encoded goal-structure." These systems' designers can no longer understand what these systems *have* "evolved into," much less anticipate into what they *will* evolve.

Large computer systems are typically not designed in the ordinary sense of the term. Though they begin with an idea—a design, if you will—which is then implemented, they soon undergo a steady process of modification, of accretion to both their control-structures and their data bases, which changes and continues to change them fundamentally. Typically too, this sort of surgery is carried out, not by the original programmers, but by people who come and go from and to other assignments. As a result, again typically, no individuals or teams of people understand the large systems. Modern large-scale systems simply *have* no authors; they have, in Minsky's words, *evolved into* whatever they have become.

Minsky long ago absolved *programmers* of responsibility for the effects of the incomprehensible systems they create precisely because their systems *are* incomprehensible:

[The] argument, based on the fact that reliable computers do only that which they are instructed to do, has a basic flaw; it does not follow that the programmer therefore has full knowledge (and therefore full responsibility and credit for) what will ensue. For certainly the programmer may set up an evolutionary system whose limitations are to him unclear and possibly incomprehensible [9].

What does it mean to *understand* a computer system at all? Minsky correctly points out that:

. . . [to] "understand" . . . implies . . . some sort of schematization—a getting at the basic principles rather than attending equally to all details however small. In that sense, "understanding" means understanding an idealized model of something rather than the thing itself! [9]

Minsky almost says that to understand something complex is to have an economical theory of the thing—and I would agree with him.

To know every line of code that constitutes a large computer program is not necessarily, not even probably, to understand the program. A theory of what the program is supposed to do is required in order to tell, for example, when the program is malfunctioning: in other words, to understand it. But this form of understanding is rendered impossible by the very way large computer systems are constructed. We are thus in precisely the situation Minsky fears: designers cannot anticipate what their machines will evolve into. And that is, as Minsky observes, a "serious problem."

Understanding the seriousness of the problem, one must surely wonder what Bell meant when he wrote: "Obviously, the information explosion can only be handled through the expansion of computerized, and subsequently automated, information systems" [5]. Perhaps a better course would be to attempt to contain the information explosion. Programmers can make a contribution by refusing to add to systems whose purposes and theories of operation cannot be explained to them.

Are there technical solutions to the problem presented by essentially incomprehensible computer systems? I don't think so. Accepting responsibility is a moral matter. It requires, above all, recognition and acceptance of one's own limitations and the limitations of one's tools. Unfortunately, the temptations to do exactly the opposite are very great. As Minsky observes, even dumb systems can be of considerable use.

On the other hand, the impressive number of comprehensible, though large, computer systems that exist in the scientific domain teach us that incomprehensibility is not *necessary* in even huge computer systems. The secret of the comprehensibility of some large computer systems in the scientific domain is that these systems are models of very robust theories. When they go wrong the errors they produce result in behaviors which contradict their theories. This should teach us that the construction of reliable computer systems in the social and political sphere awaits, not so much the results of computer science research but rather, a deeper theoretical understanding of the human condition. The limit, then, of the extent to which computers can help us cope with the world of human affairs is determined by our ability to assess our situation honestly, and our ability to know ourselves.

No discussion of the computers' role in the emergent Information Society would be complete without an appraisal of artificial intelligence (AI). This is because the spirit of artificial intelligence pervades the ethos of so much of the rest of the computer practicum. Sidney Fernbach, head of the U.S. government's Livermore Computation Laboratory, invokes an absurd vision of the potential use of artificial intelligence in science and in education:

The scientist experiences and learns to understand physical phe-

Questions that are almost never asked: Who is the beneficiary of our much-advertised technological progress and who are its victims?

nomena throughout his entire life. His most active thoughtful years are relatively few. The experience of large numbers of scientists can be put into the data banks of computer systems. *The computer then can be programmed to sort through all this information and come up with "original" ideas.* . . . Thus far we have provided for bookkeeping functions, data retrieval, problem solving in both numerical *and analytic bases, a reasoning system stocked with all the scientific knowledge in the world.* This latter system should not be restricted to science alone. Our educational facilities in general need to have the information in the Library of Congress at the fingertips of the teachers and students. This can be provided as the greatest educational tool in the world (emphasis mine) [10].

Artificial intelligence, much like real intelligence, has been extraordinarily resistant to precise definition. But there seems to be general agreement that it must be able, to use Fernbach's words, to "come up with 'original' ideas." There is also a widespread consensus that the production of original ideas has much to do with the application of analogies and metaphors. As Minsky says: ". . . in analogy lies the secret of really useful learning; a way to apply something learned in one situation to a problem in a quite different area" [10].

Minsky then discusses a program written by Thomas Evans "that proposed solutions to Geometry Analogy IQ test problems, and achieved performance resembling those of teenagers—of course, only in this restricted micro-world." Obviously, Minsky thinks this program is of very great importance to artificial intelligence.

The Evans program is given descriptions of two geometric figures A and B (the source figures) and C, D, E, and F (a small set of target figures). The problem is to select one of the figures D, E, or F such that C is to the selected figure as A is to B. A and B may be related in that, for example, some subfigure of A, A1, is *above* another subfigure of A, A2, while in B the corresponding subfigure B1 is to the *right* of B2. Given that the set of possible relationships of subfigures to one another is very small, it is possible to specify rules which govern how source figures are transformed. The program's problem then becomes to find a rule which transforms C into one of the target figures such that that rule most closely resembles the rule which transformed A1 into A2 in the original problem statement.

A metaphor is fundamentally a borrowing between and intercourse of thoughts, a transaction between contexts [11]. The extent of the creative analogical reach of a metaphor is always surprising. Its power to yield new insights depends largely on the richness of the contextual frameworks it fuses. Newton fused the contextual framework consisting of the behavior of everyday objects in the material world, for example apples falling to the ground, with that of the solar system, and produced the remarkable idea that the moon is falling to Earth.

Do the processes in Evans' program have much to do with processes used for coming up with original ideas through analogy and metaphor? This is an extraordinarily important question. For, in effect, Minsky claims that Evans' program and those that have followed the general methods it pioneered are achievements in a progression that terminates in the realization of true computer creativity.

The answer seems to me to be obvious. Truly creative thought gains its power from the combination of hitherto disparate contexts. All analogical reasoning programs artificial intelligence has produced so far are *given* the relevant criteria of similarity they need: that is, the two frameworks which are to be fused. This is not to criticize the quite clever programs produced to date; it is rather to illustrate on what profoundly and fundamentally misguided bases some of the most crucial concepts of artificial intelligence are built.

The modern computer was born from the womb of the military. As with so much other modern technology of the same parentage, almost every technological advance in the computer field has had its residual payoff—fallout—in the civilian sector. Still, computers were first constructed in order to calculate efficiently how best to drop artillery shells on people. Probably the largest fraction of computers devoted to a single purpose today are still dedicated to cheaper, more accurate ways of killing ever larger numbers of human beings.

What then *can* we expect from this strange fruit of the human genius?

We can expect the kind of euphoric forecasting and assessment which fills the popular and some of the scientific literature. Unrelated to computers per se, this is a characteristically American tradition of thought. We have seen many other examples of it—and these may be instructive.

Americans thought that universal schooling would uplift the masses and ensure a happy, prosperous, democratically-governed society. While almost all American youngsters today are forced to attend school during the whole of adolescence, our primary and secondary schools have not become centers of learning, or even centers where elementary reading and writing *can* be taught at all. Government reports document that America's young people are largely functionally illiterate. As a university professor, I can testify that not many youngsters recruited from among the best and the brightest can compose a single paragraph of standard English prose.

We can also expect that the computer scientists and other scholars

Will our children be able to live with the world we are here and now constructing?

who claim to have made themselves authorities in this area will, on the whole, see the nude emperor's magnificent garments more clearly than any one else. Some of us will find their accounts unrealistic because they are plainly silly. For example, the distinguished Princeton Professor of Public and International Affairs, Robert Gilpin, writes:

> In order to exercise economic power, a nation must be able to process vast amounts of data. The classic case in point is the Arab petroleum boycott against the West following the October 1973 Arab-Israeli War. Without sophisticated data processing capabilities, the Arab oil producers *could not have* kept track of Western oil tankers, refinery output, etc., and thereby enforced their embargo. Moreover, given the complexity of the oil industry and the potentialities of cheating by cartel members, it is doubtful if the Organization of Petroleum Exporting Countries (OPEC) would remain intact without the benefit of electronic data processing (emphasis mine) [12].

Oil tankers spend weeks at sea. An old-fashioned clerk with a quill pen could keep track of them on the back of a few large envelopes. And there have been effective cartels since at least the rise of modern capitalism.

We need not credit computers for accomplishments with which they have nothing to do. They can be realistically credited with having made possible some easing of the lives of some people. Modern airline reservation systems, for example, have made it easier for me to travel. Computers have radically transformed many aspects of astronomy, and without computers space flight would have been impossible. The computer has done some good.

But some questions are almost never asked:

- Who is the beneficiary of our much-advertised technological progress and who are its victims?

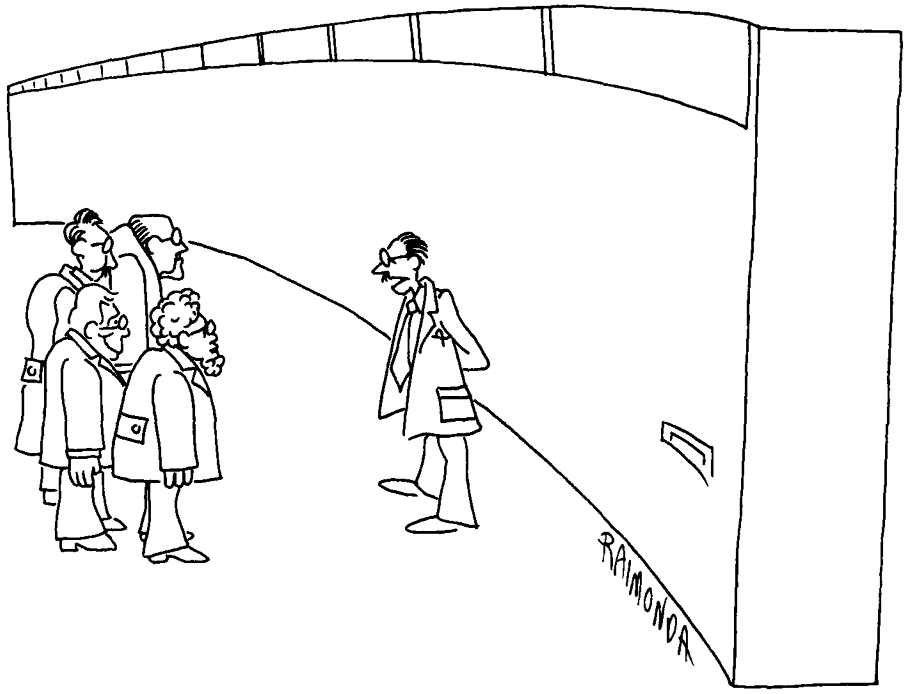

"What I want to know, gentlemen, is, 'Who taught the Multi-Unit 3000 to lie?' "

- What *limits* ought we, the people generally and scientists and engineers particularly, to impose on the application of computation to human affairs?
- What is the impact of the computer, not only on the economies of the world or on the war potential of nations, etc., but on the self-image of human beings and on human dignity?
- What irreversible forces is our worship of high technology, symbolized most starkly by the computer, bringing into play?
- Will our children be able to live with the world we are here and now constructing?

Much depends on answers to these questions.

Notes

1. M. L. Dertouzos, "History as Prologue: The Road to Science and Utopia," in *The Future Study on the Impact of Computers and Information Processing,* Dertouzos, ed. (Cambridge, Mass: The M.I.T. Press, forthcoming).
2. S. Papert, "Computers and Education," in Dertouzos [1].
3. J. Moses, "The Computer in the Home," in Dertouzos [1].
4. B. O. Evans, "Computers and Communication," in Dertouzos [1].
5. Daniel Bell, "The Social Framework of the Information Society: The Coming Age of Communication and Control," in Dertouzos [1].
6. Jay Forrester, testimony before House Committee on Banking and Currency, Subcommittee on Urban Growth, Oct. 7, 1970 (Washington, D.C.: GPO), part III, pp. 205-265.
7. A. J. Perlis, "Current Research Frontiers in Computer Science," in Dertouzos [1].
8. M. Minsky, "Computer Science and the Representation of Knowledge," in Dertouzos [1].
9. M. Minsky, "Steps Toward Artificial Intelligence," in *Computers and Thought,* Edward A. Feigenbaum and Julian Feldman, eds. (New York: McGraw-Hill, 1963), p. 447.
10. S. Fernbach, "Scientific Use of Computers," in Dertouzos [1].
11. I. A. Richards, *The Philosophy of Rhetoric* (Oxford, England: Oxford University Press, 1936), p. 93.
12. R. Gilpin, "The Computer and World Affairs," in Dertouzos [1].

An expanded version of this article is to appear in a volume on the future impact of computers and information processing to be published by The M.I.T. Press in the spring of 1979.