ATTITUDES TOWARD INTELLIGENT MACHINES

by Paul Armer

"A BIRD IS AN INSTRUMENT WORKING ACCORDING TO MATHEMATICAL LAW, WHICH INSTRUMENT IT IS WITHIN THE CAPACITY OF MAN TO REPRODUCE WITH ALL ITS MOVEMENTS."
Leonardo da Vinci (1452–1519)

This is an attempt to analyze attitudes and arguments brought forth by questions like "Can machines think?" and "Can machines exhibit intelligence?" Its purpose is to improve the climate which surrounds research in the field of machine or artificial intelligence. Its goal is not to convince those who answer the above questions negatively that they are wrong (although an attempt will be made to refute some of the negative arguments) but that they should be tolerant of research investigating these questions. The negative attitudes existent today tend to inhibit such research (MacGowan, 1960).[1]

History

Before examining the current arguments and attitudes toward artificial intelligence, let us look at some of the history of this discussion, for these questions have been around for a long time.

Samuel Butler (1835–1902), in *Erewhon and Erewhon Revisited* (1933), concocted a civil war between the "machinists" and the "anti-machinists." (Victory, incidentally, went to the "anti-machinists.") Butler stated "there is no security against the ultimate development of mechanical consciousness in the fact of machines possessing little consciousness now" and speculated that the time might come when "man shall become to the

[1] Almost an entire book, *Computers and Common Sense, The Myth of Thinking Machines*, has been devoted to condemning artificial intelligence research (Taube, 1961). Readers who have been exposed to this book should refer to reviews of it by Richard Laing (1962) and Walter R. Reitman (1962), particularly the former.

machines what the horse and dog are to us." Discussion of this topic apparently took place in Babbage's time (1792–1871), for the Countess of Lovelace commented on it, negatively, in her writings on Babbage's efforts (Bowden, 1953). The topic came into prominence in the late 1940's when Babbage's dreams became a reality with the completion of the first large digital computers. When the popular press applied the term "giant brains" to these machines, computer builders and users, myself included, immediately arose to the defense of the human intellect. We hastened to proclaim that computers did not "think"; they only did arithmetic quite rapidly.

A. M. Turing, who earlier had written one of the most important papers in the computer field on the universality of machines (1936, 1937), published in 1950 a paper entitled, "Computing Machinery and Intelligence." In it he circumvented the problem of properly defining the words "machine" and "thinking" and examined instead the question of a game wherein an interrogator, who can communicate with a human and a machine via teletype, but does not know which is which, is to decide which is the machine. This is now known throughout the computer field as "Turing's Test."

Discussion of machine intelligence died down (but not out) in the early and mid-1950's but has come back in the last several years stronger than ever before. In fact, it has recently invaded the pages of *Science* (MacGowan, 1960; Wiener, 1960; Taube, 1960; Samuel, 1960*b*).

A Way of Thinking about Thinking

Before beginning an examination of the negative arguments, allow me to introduce a concept which will aid in discussing these arguments and which may help resolve some of the semantic difficulties associated with discussions of "Can machines think?" Like Turing, I avoid defining "to think." Instead, observe that thinking is a continuum, an *n*-dimensional continuum. This notion is certainly not new, for it has existed since man first compared his mental abilities with another man's, and it is implicit in all of the positive arguments on machine intelligence. Psychologists long ago developed "intelligence quotient" as a yardstick in this continuum, and their concept of "factors" is indicative of the *n*-dimensionality of the continuum of intelligence. The use of the one-dimensional "I.Q." is obviously an oversimplification of reality. Although the concept of an *n*-dimensional continuum for intelligence is not new, and although it is implicit in many discussions of artificial intelligence, it is rarely stated explicitly.

An analogy may be drawn with the continuum of the ability to transport. With respect to speed in transporting people from New York to Los Angeles, the jet airplane of today outshines all other existing transportation vehicles. But it does not compare favorably, costwise, with ships for transporting newsprint from British Columbia to California. Existing commercial jet transports cannot transport people from one lake to another. A Cadillac may be the most comfortable vehicle to transport people short distances over a good network of roads, but it is hardly a substitute for the jeep in the environment of ground warfare—the jeep's forte is versatility and flexibility. In this dimension, in the continuum of the ability to transport, man outshines the jeep, for man can go where jeeps cannot, just as the jeep can go where Cadillacs cannot. But men cannot carry the load that a jeep can nor can men move with the speed of the jeep.

Similarly, comparisons can be made between men and machines in the continuum of thinking. If there is objection to the use of the word "thinking," then "ability to process information" or some similar term can be used. But it must be admitted that there exists some continuum of behavior in which men and machines coexist and in which they can be compared. (See Fig. 1.)

An n-dimensional continuum is difficult to draw when n is large, so let's examine a two-dimensional one, realizing that reality is far from being that simple. With respect to raw speed, machines outdo men, but when it comes to the sophistication of the information processes available, machines look pretty poor. This dimension deserves further discussion. While the repertoire of today's machines is quite simple—a few basic arithmetic operations and comparisons—man's information processes are very complex. Let me illustrate this point with the following incident. We have all had the experience of trying to recall the name of a person we have once met. On a particular occasion Dr. Willis Ware and I were both trying to recall an individual's name. We recounted to one another his physical characteristics, where he worked, what he did, etc. But his name eluded us. After some time, I turned to Dr. Ware and said, "His name begins with a 'Z'." At which point he snapped his fingers and correctly said, "That's it, it's Frizell!"

Now, of course, the basic question is "Can the machines' capabilities in this dimension be improved?" Let me turn the question around—Is there any evidence that they cannot? I know of none. In fact, over the last decade I think impressive progress has been made. It's easy to underestimate the advances, for "intelligence" is a slippery concept. As Marvin Minsky put it, "You regard an action as intelligent until you understand it. In explaining, you explain away" (1959a).

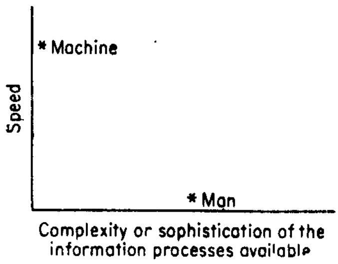

Figure 1

Today's computers, even with their

limited capability in the sophistication dimension, have had tremendous impact on science and technology. Accomplishments of the last decade in the fields of nuclear energy, missiles and space would have been impossible without computers. If we can push the capabilities of computers[2] further out in the sophistication dimension, won't they have an even greater impact? *In this context then, the goal of research on artificial intelligence can be stated—it is simply an attempt to push machine behavior further out into this continuum.*

It is irrelevant whether or not there may exist some upper bound above which machines cannot go in this continuum. Even if such a boundary exists, there is no evidence that it is located close to the position occupied by today's machines. Is it not possible that we might one day understand the logical processes which went on in Dr. Ware's head and then mechanize them on a machine? We obviously will not achieve such a goal unless someone believes that it is possible and tries to do it. One does not have to believe that the boundary is nonexistent in order to try; one need only believe that the boundary is much further out than the position occupied by today's machines.

Intelligent Machines and Today's Digital Computer

A common attitude toward today's computers is that such machines are strictly arithmetic devices. While it is true that machines were first built to carry out repetitive arithmetic operations, they are capable of other, nonnumeric tasks. The essence of the computer is the manipulation of symbols—it is only a historical accident that the first application involved numeric symbols. This incorrect notion of the computer as a strictly numeric device results in the inability of many to conceive of the computer as a device exhibiting intelligent behavior, since this would require that the process be reduced to a numerical one. The reaction of many people to statements about intelligent behavior by machines seems to indicate that they take such statements to imply complete functional equivalence between the machine and the human brain. Since this complete functional equivalence does not exist, such people believe they have thereby debunked intelligent machines. Their argument is hollow since this equivalence was never implied. Intelligent behavior on the part of a machine no more implies complete functional equivalence between machine and brain than flying by an airplane implies complete functional equivalence between plane and bird.

The concept of comparing the behavior of men and machines in an *n*-dimensional continuum recognizes differences as well as similarities.

[2] I make no distinction here between the attributes of the computer and those of the program which controls the computer.

For example, a common argument against machine intelligence is that the brain is a living thing—the machine is not. In our continuum we simply recognize the dimension of living and note that machines and men occupy different positions in this dimension.

While I do believe that today's digital computers can exhibit intelligent behavior, I do not hold that the intelligent machines of the 1970's will necessarily resemble today's machines, either functionally or physically. In particular, in my desire to see machines pushed further out in the continuum of intelligence, my interests in the dimension of speed are very minor; the organizational aspects (sophistication of the information processes) are obviously much more important. Likewise, I hold no brief for the strictly digital approach; a combination of analog and digital equipment may prove to be better. I do not mean to disown the digital computer, for it will be a most important tool in the endeavor to advance in our continuum.

Some of the Negative Arguments

1. The Argument of Invidious Comparison

Considering the behavior of men and machines in the context of intelligence being a multidimensional continuum, an argument that a machine cannot play chess because "it could only operate on standard-size pieces and could not recognize as chessmen the innumerable pieces of different design which the human player recognizes and moves around quite simply" (Taube, 1960) is like saying that the Wright brothers' airplane could not fly because it could not fly nonstop from Los Angeles to New York nor could it land in a tree like a bird. Why must the test of intelligence be that the machine achieve identically the same point in the continuum as man? Is the test of flying the achievement of the same point in the continuum of flying as that reached by a bird?

2. The Argument of Superexcellence

Many of the negativists[3] seem to say that the only evidence of machine intelligence they will accept is an achievement in our continuum seldom achieved by man. For example, they belittle efforts at musical composition by machine because the present output compares miserably with that of Mozart or Chopin. How many *men* can produce music that compares favorably? The ultimate argument of this kind occurred at a recent meeting in England, during which a discussant stated that he would not accept the fact that machines could think until one proved the famous

[3] The terms "negativists" and "positivists" are used in this report to classify those who do not and those who do, respectively, believe machines can exhibit intelligent behavior. Of course, variations of degree exist.

conjecture of Fermat, better known as Fermat's last theorem. By this logic one concludes that, to this date, no man has been capable of thinking, since the conjecture remains unproven.

3. The Argument by Definition

There are many variations of this type of argument. For example, some negativists want to include in their definition of intelligent behavior the requirement that it be carried out by a living organism. With such a definition, machines do not behave intelligently. However, there does still exist machine behavior which can be compared with human behavior. To conclude that research on the simulation of such human behavior with a machine is wrong, as some have done, because the machine is not living, is like concluding that research on the simulation of the functions of the human heart with an artificial heart is wrong because the artificial organ is not a living one.

4. The Argument by Stipulation

An examination of the arguments advanced by the negativists reveals that many of them are not arguments at all, but only statements. They dismiss the notion out of hand, saying things like, "Let's settle this once and for all, machines cannot think!" or "A computer is not a giant brain, in spite of what some of the Sunday supplements and science fiction writers would have you believe. It is a remarkably fast and phenomenally accurate moron" (Andree, 1958).

5. The Argument by False Attribution

Typical of this type of argument is the following:

> *The Manchester machine which was set to solve chess problems presumably proceeded by this method, namely by reviewing all the possible consequences of all possible moves. This, incidentally, reveals all the strength and weakness of the mechanism. It can review far more numerous possibilities in a given time than can a human being, but it has to review all possibilities. The human player can view the board as a whole and intuitively reject a number of possibilities. The machine cannot do either of these* (Hugh-Jones, 1956).

The statements about machine behavior in the above quotation are simply not true. While it is true that some of the early approaches to chess-playing machines were in the nature of attempts to review *all* possibilities in limited depth (Kister *et al.*, 1957), this is not the only way in which the problem can be approached. The chess-playing routine of Newell, Shaw, and Simon (1958*b*) does *not* examine all possibilities. And those which it does consider it examines in varying detail. The routine rejects moves which appear to be worthless; it selects moves which appear to be good ones and examines them in depth to ascertain that they are indeed good. An earlier routine developed by this same team to prove theorems in logic (Newell, Shaw and Simon, 1957*a*) did not examine all possible proofs—to do so with today's computers would literally take endless time. Rather, the routine searched through the maze of possible proofs for ones which looked promising and investigated them. It relied on knowing which approaches had worked before. Most of those who scoff about research on artificial intelligence turn out to be unaware of the details of what is going on in such research today; it is little wonder that they frequently make erroneous statements about the field.

6. The Argument by False Extrapolation

This class of argument is typified by extrapolations based on assumptions that machine properties are invariant. For example:

> *The human memory is a filing system that has a far greater capacity than that of the largest thinking machine built. A mechanical brain that had as many tubes or relays as the human brain has nerve cells (some ten billion) would not fit into the Empire State Building, and would require the entire output of Niagara Falls to supply the power and the Niagara River to cool it. Moreover, such a computer could operate but a fraction of a second at a time before several thousand of its tubes would fail and have to be replaced* (Troll, 1954).

The point is tied to the vacuum tube (the article was written in 1954) and has therefore already been weakened by the appearance of the transistor, which requires less space and power and is considerably more reliable than the vacuum tube. An offsetting development is that the estimate of the number of nerve cells is undoubtedly too low. However, on the horizon are construction techniques involving the use of evaporated films, where the details of the machine will not be visible under an optical microscope (Shoulders, 1960). It seems reasonable to expect that it will be possible with these techniques to house in one cubic foot of space the same number of logical elements as exist in the human brain. Power requirements will be trivial.

7. The Obedient Slave Argument

One often hears statements like "The machine can only do what it is told to do." People who advance this obedient slave argument would seem to be thinking that they are countering others who have pointed to a large conglomeration of unconnected transistors, resistors and electronic components, and said "It thinks." Certainly man is involved in machine intelligence—so are parents and teachers in human intelligence. Do we deny flying to an airplane because a man is piloting it or even to an unmanned flight because a man designed it?

The negativists who say "the machine can only do what it is told to do" overlook the fact that they have not qualified their statement as to what is the limit of what the machine can be told to do. What evidence exists concerning the location of that limit? Might it not become possible to tell a machine to learn to do a given task, a task usually considered to require intelligence? Many of the tasks being accomplished with computers today were not considered possible ten years ago.

Recent Computer Tasks and Milestones

The mounting list of tasks which can now be carried out on a computer but which we normally consider requiring intelligence when performed by humans, includes such things as:

Proving theorems in logic and plane geometry (Newell, Shaw and Simon, 1957*a*; Gelernter, 1960*a*)

Playing checkers and chess (Samuel, 1959; Newell, Shaw and Simon, 1958*b*)

Assembly line balancing (Tonge, 1961*a*)

Composing music (Hiller and Isaacson, 1959)

Designing motors (Goodwin, 1958)

Recognition of manual Morse code (Gold, 1959)

Solving calculus problems (Slagle, 1961)

The collection of capabilities which have been ascribed solely to humans in the past is being slowly chipped away by the application of computers. Space precludes going further into the evidence for machine intelligence; this topic is well covered in the articles previously cited and in other papers (Newell, Shaw and Simon, 1956; Milligan, 1959; Minsky, 1961*a*). Such evidence is, of course, the basis for many of the arguments advanced by the positivists.

To prove that machines *today* do *not* exhibit intelligence, it is only necessary to define a lower bound in our continuum which is above the behavior exhibited by the machines and then say that behavior above that bound is intelligent and below it is not intelligent. This is a variant of the proof by definition. Many who use this gambit have been redefining the lower bound so that it is continually above what machines can do today. For example, we find

> *Perhaps the most flexible concept is that any mental process which can be adequately reproduced by automatic systems is not thinking* (Meszar, 1953).

This redefinition may not be done consciously. A skill which seems highly intelligent in others becomes much less impressive to us when we acquire that skill ourselves. It would be useful to have at hand some milestones for the future. Turing's test is one such milestone (1950) but additional ones are needed. To this end a clearly defined task is required which is, at present, in the exclusive domain of humans (and therefore incontestably "thinking") but which may eventually yield to accomplishment by machines.

Rivalry Between Man and Machines

There is a strong personal factor in the attitude of many negativists. I'm sure it was a major factor in my being a negativist ten years ago. To concede that machines can exhibit intelligence is to admit that man has a rival in an area previously held to be within the sole province of man. To illustrate this point, let me quote from a letter received at RAND:

> *. . . semantics may have a lot to do with the degree of enthusiasm for supporting research in this area (artificial intelligence). Subjectively, the terms "intelligent machine" or "thinking machine" disturb me and even seem a bit threatening: I am a human being, and therefore "intelligent" and these inhuman devices are going to compete with me and may even beat me out. On the other hand, if the very same black boxes were labelled "problem solver," or even "adaptive problem solver," they would seem much more friendly, capable of helping me in the most effective way to do things that I want to do better, but, best of all, I'd still be the boss. This observation is wholly subjective and emotional. . . .*

Another explanation of why some negativists feel the way they do is related to what might be called the "sins of the positivists." Exaggerated claims of accomplishments, particularly from the publicity departments of computer manufacturers, have resulted in such a strong reaction within the scientific community that many swing too far in the opposite direction.

Da Vinci and Flying

At this point allow me to paraphrase the quotation of da Vinci's, with which this paper was begun, and also, with the benefit of hindsight, expand on it somewhat. Thus, he might have said:

> *When men understood the natural laws which govern the flight of a bird, man will be able to build a flying machine.*

While it is true that man wasted a good deal of time and effort trying to build a flying machine that flapped its wings like a bird, the important point is that it was the understanding of the law of aerodynamic lift (even though the understanding was quite imperfect at first) over an airfoil which enabled men to build flying machines. A bird isn't sustained in the air by the hand of God—natural laws govern its flight. Similarly, natural laws govern what went on in Dr. Ware's head when he produced "Frizell" from my erroneous but related clue. Thus, I see no reason why we won't be able to duplicate in hardware the very powerful processes of association which the human brain has, once we understand them. And if man gained an understanding of the processes of aerodynamics, may he not also obtain an understanding of the information processes of the human brain?

There are other facets to this analogy with flight; it, too, is a continuum, and some once thought that the speed of sound represented a boundary beyond which flight was impossible.

Approaches to the Problem of Building an Intelligent Machine

This topic can perhaps be expounded best with another analogy. Suppose we are given a device which we know exhibits intelligent behavior because we have observed it in action. We would like to build a machine which approaches it in capability (or better yet, exceeds it). We bring in a group of men to study the basic components of the device to understand how they work. These men apply pulses to subsets of the leads, and observe what each component does; they try to understand why the device behaves as it does in terms of basic physics and chemistry. They also seek to learn how these components function in subassemblies.

A second group of men approach the problem from the point of view that the device is a "black box" which they are not able to open. This group observes that some of the appendages of the device are obviously input devices while others are output devices. They observe the device in operation and attempt to theorize how it works. They proceed on the basis that it will not be necessary that the machine they are to construct have the same basic components as exist in the device under study. They believe that if they can understand the logical operation of the existing device, they can duplicate its logic in their own machine, using components they understand and can make.

This second group makes conjectures about the logical construction of the device and tries these conjectures out on a computer which they have at hand. These theories are very crude at first and do not mirror the behavior of the "black box" very well, but over time the resemblance improves.

Because we learned a lesson from the effort spent on attempting to build a flying machine that flapped its wings, we set a third group to work studying "intelligence and information processing" per se and building up a science in the area.

There is much common ground among the three groups and they keep each other posted on results to date. Furthermore, they all use computers to aid them in their research. The groups combine their know-how along the way to build better computers (low-I.Q. intelligent machines) on which to try out their conjectures. Eventually, the three groups "come together in the middle" and build a machine which is almost as capable as our model. They then turn to the task of building an even better one.

In the real-life situation of studying the human brain, the first group, studying components and assemblies thereof, is represented by physiological work. The second, or "black-box" group, is represented by psychological efforts to explain human mental activity. This analogy represents, I believe, a plausible scenario for the way things might go in trying to understand the human mind.

Russian Attitudes

Our examination thus far has been Western in origin; in view of the impact that achievement of the goals of research on artificial intelligence would have on the technological posture of the United States vis-à-vis the Soviet Union, it might be interesting to look at Soviet attitudes toward intelligent machines. As one might suspect, Soviet attitudes have been quite similar to Western ones. Positivists and negativists exist, and each camp advances the same sort of arguments as their Western counterparts. For example, there are negativists who advance the obedient slave argument. Academician S. A. Lebedev, head of the Institute of Precise Mechanics and Computational Techniques and host to the U.S. Exchange Delegation in Computers which visited the USSR in the last two weeks of May, 1959 (of which I was a member), on two occasions dismissed my questions concerning his attitude toward intelligent machines with the statement "Machines can do no more than they are instructed to do."

Their literature is filled with discussions of comparisons between men and machines. In 1961, an entire book, *Philosophical Problems of Cybernetics* (1961), was published on this topic. It was obvious from the questions asked of our delegation by the Russians about Western attitudes that it is a hotly debated issue. In the USSR, research on artificial intelligence is a part of cybernetics, the term coined by Wiener (1948) and now a household word in the Soviet Union. Cybernetics is also used as an umbrella term for research in automatic control, automation, computers, programming, information retrieval, language translation, etc. It is universally recognized as an area related to both men and machines, and the requirement for an interdisciplinary (engineering, mathematics, computing, biology, psychology, physiology, physics, chemistry, linguistics, etc.) approach to such research is also recognized.

As in the West, the use of the term "giant brains" in the late 1940's resulted in a massive revulsion among the Soviet scientific community, and universal rush to the defense of the human mind. The degree of the revulsion was such that several Soviet writers have blamed it for the fact that Russia presently lags the U.S. in the digital-computer field (Shaginyan, 1959). One finds frequent references in the Russian literature to the existence of a negative attitude toward cybernetics, and to the persistence of this attitude for a period of about ten years.

Soviet literature on cybernetics frequently gives credit to Wiener, von Neumann, and other Westerners for pioneering the field. It also contains many references to the work of Pavlov and mixes in much political discussion of communism vs. capitalism, and even of Marx and Lenin. For example, we have:

> *Karl Marx was the first to make use implicitly and anticipatingly of cybernetical ways of thought, or to express it more pointedly, Karl Marx was the first cybernetician!* . . . (Klaus, 1960).

There are some strong positivists in the USSR. For example, I. A. Poletayev has stated "nothing except prejudice and superstition allow one to deny with assurance today the possibility that the machine will pass, in the end, that limit beyond which consciousness begins" (1958). Other strong positivists include S. L. Sobolev (an Academician and a well-known mathematician) and A. A. Lyapunov (1960). We also find:

> *. . . Thus, the perfecting of computer machines involuntarily leads us to the need to create a model of the brain. . . . Also, one of the most effective methods of studying intra-cerebral processes involve experiments carried out in electrical models of the brain. . . .*
> *But cybernetics has its critics too. These are skeptics. One can find them among scientists and among ordinary citizens, at times also among administrative personnel. These skeptics reject this branch of science and deny it the right of existence. . . . In rejecting this science, they generally state that the very thought of comparing a machine to a human being is an insult* (Moiseyev, 1960).

The majority of Soviet workers appear to recognize (implicitly, at least) the continuum discussed in this report, and argue that while there does exist an upper bound above which machines cannot go, it is not possible to determine the location of that bound. For example:

> *As a result we arrive at the conclusion that a machine can perform all the intellectual human functions which can be formalized . . . But what can be formalized? . . . Upon brief reflection we con-*

clude that it is impossible in principle to answer this question (Kolman, 1960).

Where Do the Russians Stand?

First of all, let us look at what they are doing in those disciplines upon which research in artificial intelligence depends: computing devices, mathematics, psychology, and physiology. With respect to computers, I can speak with firsthand knowledge, for, as mentioned earlier, I spent two weeks in 1959 visiting Soviet computer installations. In my opinion, they are somewhat behind us in the actual construction of machines, particularly with respect to input/output equipment and to numbers of machines (Ware, 1960; Feigenbaum, 1961*c*). However, there is nothing fundamentally lacking in their state of the art. The quantity of machines is not as important to research as an offhand comparison of numbers of machines might indicate, since none of their machines is devoted to such things as social-security records, subscription fulfillment, or airline reservations. In assessing a comparison of this kind, one always wonders how much of the iceberg we do not see. When visiting the IBM plant in California, Khrushchev said about computers, ". . . for the time being we're keeping them a secret."

The Russians started work on computers after we did, but they have certainly narrowed the gap. Furthermore, they are giving high priority to the computing field. In their announcement concerning the decentralization of responsibility for research, an exception was made for computers, along with fusion, space activities, high-temperature metallurgical research, and certain areas of chemistry; these research areas remained centralized under the cognizance of the Academy of Sciences. Of course, the Russians are interested in spurring the computer field for reasons other than intelligent machine research. There is no reason to believe that future Russian research on intelligent machines need be hampered by the computer tools available to them, although machine time is in short supply today.

In mathematics the Russians have had an outstanding reputation for many decades. In computer mathematics I have no doubts that, in general, they excel the West. One of the things which impressed our delegation, and other delegations before ours (Carr *et al.,* 1959), was the number of outstanding mathematicians now working in the computer field. Unfortunately, many U.S. mathematicians view computers as a glorified slide rule of interest only to engineers, or as an expensive sorting device of interest to businessmen with clerical problems.

Since psychological research on mental processes and neurophysiological research on structure and activity of the brain both play a vital suggestive role in the attempt to construct intelligent machines, progress by the Soviets in these disciplines is of considerable interest. Although psychology was severely inhibited during the Stalin era, a renaissance of impressive proportions has taken place within the last decade. Physiology, less inhibited in the previous era, is in even better shape. The best available evidence indicates that Russian neurophysiology is dynamic, innovative, and up to date. The researchers are competent and generally sophisticated; their laboratories are modern and well equipped.

The Soviets have demonstrated a knack for focusing talent and resources on important applied problems. I believe that the Soviets regard artificial intelligence as one such problem area, and that the best of modern Soviet psychology and neurophysiology will be recruited into the search for solutions. With respect to physiological research, the following is of interest:

> *Essentially, we (the Western World) have not found the physiochemical principles of neural activity, whereas the Russians have not seriously sought them. However, the current 7-year plan for physiology as presented in a recent editorial by D. A. Biriukev in the Sechenov Physiological Journal of the USSR calls for precisely this goal* (Freeman, 1960).

A recent visitor of the USSR reports that Soviet physiologists appear to be under pressure to produce explanations for human behavior which can be incorporated into machines. He further reports that their work is apparently under security wraps.

Russian Emphasis on Artificial Intelligence Research

I went to the Soviet Union convinced they were putting a great deal of emphasis on research in artificial intelligence. Possibly this predisposition influenced what I thought I saw. I also want to emphasize that *I was impressed, not by any substantive results, but by their apparent conviction that this was an important research area.*

In one institute, in response to my question about the problem of simulating the brain with a computer, I was told "It is considered *the* number one problem." The emphasis on "the" was the speaker's; the statement was made in English. At another institute, when Professor L. I. Gutenmacher, head of the Laboratory for Electrical Modeling, told us that the charter of his laboratory was the modeling of human mental processes, I asked him if he had difficulty obtaining financial support for such exotic research. His response was "No, not at all; the President of the Academy of Sciences is convinced that this is an important field for research." There is evidence that he has been given ample support. I was told that his laboratory, which was formerly (and still is ostensibly) a part of the Institute of Scientific Information, had all the status of an institute, being separately funded and reporting directly to the Presidium of the Academy of Sciences. Gutenmacher's laboratory is apparently responsible for mechanizing the functions of the Institute of Scientific Information, which is a large, centralized, information retrieval system for scientific information from all over the world.

Despite much effort, our delegation was unable to visit Gutenmacher's laboratory. To my knowledge, no Westerner has done so; in fact none had met Gutenmacher before our delegation. Some in the U.S. have concluded from this denial of entry to his laboratory that there was nothing to be shown. However, its work may be classified, as Khrushchev indicated. But whether or not anything is being accomplished is not pertinent to the point that the President of the Soviet Academy of Sciences, a man with much power and resources, believes that modeling human mental activities is possible, that he recognizes the importance of research in this field, and that he is devoting considerable resources to this end.

What are some of the other indications about Soviet attitudes toward research on intelligent machines? As previously mentioned, cybernetics is a household word in Russia. Much is being written on the subject, in journals and in the popular press. There appears to be an effort in the popular writings to legitimatize such research as being in harmony with communism. For example, recall the earlier quote about Marx (Klaus, 1960).

With respect to professional writing on machine intelligence, a journal entitled *Problems of Cybernetics* was started in 1958; seven hard-cover volumes have appeared to date (Lyapunov, 1960, 1961). Since 1955, seminars on cybernetics have been held at the University of Moscow. These seminars are aimed at bringing together scientists from various disclipines. Similarly, the editors of *Problems of Cybernetics* state that their aim "is the unification of the scientific interests of those working in different fields of science concerned with cybernetics."

There seems to be widespread recognition for the necessity of an interdisciplinary approach to problems of cybernetics. Article after article appeals to personnel from the various disciplines to get together. How much effect these appeals and seminars have is unknown. During our visit to the Soviet Union, we were told that some 500 physicists had been transferred to the biological sciences. We talked with I. M. Gelfand, a world-famous mathematician now working in the physiological field. He began studying the brain but switched to the heart, which he believes to be much simpler. With knowledge gained from studying the heart, he will return to the study of the brain. We were also told that other mathematicians were working on psychological and physiological problems.

Within the Soviet Academy of Sciences, there exists a "Scientific Council on Cybernetics." This council is headed by A. I. Berg and apparently reports directly to the Presidium of the Academy (Berg, 1960). To my knowledge, there is no evidence of any effect this council may be having in coordinating, controlling, or encouraging research in cybernetics. Outside of Moscow, individual researchers appear to operate entirely on their own, with little communication with other such researchers, and with only meager support. However, one does occasionally encounter references to the formation of new groups and laboratories for such work.

There is some evidence that machine time (until recently in critically short supply) has been made available for work in this area. *Moscow News* of August 12, 1961, has an article on musical composition and medical diagnosis on a computer while the issue of September 2, 1961, discusses chess playing by machines and the deciphering of ancient Mayan manuscripts.

In closing this topic, a quotation which appeared in the February, 1959, issue of *Fortune* is pertinent. Frank Pace, Jr., then president of General Dynamics Corporation, in warning us not to overlook nor be surprised by Russia's capacity to concentrate in specific areas, said:

> *If the area has real military or psychological value to them, they'll put massive concentration on it, and achieve results all out of proportion to the general level of their technical ability.*

The Importance of Research in Artificial Intelligence

I have indicated my feeling that research aimed at pushing machines further out in the continuum of intelligence is very important. Today's computers are helping advance the frontiers of man's knowledge in many fields; computers now pervade almost all scientific disciplines. (The fact that they pervade the field of research on intelligent machines means that such research will feed on itself.) The use of computers in research has been a key factor in the explosion of knowledge we have witnessed in the last decade. Their contribution to date has stemmed largely from their speed in doing arithmetic and the reliability with which they do it. As we move out in the continuum of possibilities, new dimensions and contributions will become important. A machine which retrieves information from a large store by complex associative processes like those inherent in Willis Ware's output of "Frizell," but which exceeds Dr. Ware in speed, reliability, and memory capacity, would be crucial in aiding scientists to cope with the flood of research results presently inundating science.

The large amount of money spent on machines today is evidence of the value placed on the computers' abilities along the dimensions of speed and reliability. If the machine's capabilities can be extended in additional dimensions, would it not be of great importance? Suppose that the boundary (if it exists at all) beyond which machines cannot go lies fairly close to the human brain in the dimension related to the sophistication of the information processing techniques used. Since it is known that the machine can exceed the human in speed and reliability, and probably in amount of memory, such a machine would approach the status of being "super-human." Of course, this is speculation; the boundary may be much lower.

We have been examining the question of the technological importance of research in artificial intelligence in the context of advancing the frontiers of knowledge for the sake of technological and scientific advancement. In such a context, there is little cause for any concern or action; progress in the field is being made at a fairly rapid pace in this country. However, since we are engaged in a technological race with the USSR, action becomes important, particularly since, in my opinion, the Russians appear to be putting much more emphasis on research in artificial intelligence than we are. Even if the Russians were not competing in this particular event of the "technological Olympics," it is an event well worth the running in that we will learn more about man and in that better machines will contribute to advancing the frontiers of knowledge in almost every discipline.

Timing

Before closing, a comment on the question "when?" It is one thing to say it is possible to push machine capabilities way out in the continuum of intelligence, but it is another thing to say when. It was over four hundred years from da Vinci to the Wright brothers. But the sands of time in the scientific world have been flowing much more rapidly of late. Advances now made in a decade compare with earlier steps which took a century. Few would have believed in 1950 that man would hit the moon with a rocket within ten years. Gutenmacher, when told recently of the Simon and Newell prediction that a machine would be chess champion within ten years (Newell and Simon, 1958*d*) said that he thought the prediction conservative; it would happen sooner.

Conclusion

It is hoped that the definition of research on artificial intelligence as an effort to push machines further out in the continuum of intelligent behavior will reduce some of the semantic difficulties surrounding discussions of such research. I feel that such research is very important to our country and that we must expand our efforts therein. To do so implies that more researchers from the related disciplines are needed. The success of our efforts will depend on how well we do in bringing the various disciplines together and on the number of well-qualified scientists who are attracted to this research area.