THE GREAT IDEAS ONLINE

May '20 *Philosophy is Everybody's Business* № 1040

INTELLECT: MIND OVER MATTER

Mortimer Adler

PART FIVE: ARTIFICIAL INTELLIGENCE AND THE HUMAN INTELLECT

5

Currently many behavioral scientists accept the prediction that computer technology will make it possible to construct machines with artificial intelligence that will enable them to do everything that human beings can do. Their performance will be indistinguishable from that of the human intellect. Underlying this prediction is the materialistic dogma that denies the immateriality of the human intellect—the dogma that supports the view that the brain is the necessary and sufficient condition of all mental acts and processes.

These projected artificial intelligence machines will not be alive, will have no vegetative powers, and, while they may have something like the sensory powers of living organisms, especially perceptive powers, they will not have consciousness or experience feelings of pleasure and pain or the emotions of anger and fear.

If all the apparent differences in kind between human and animal behavior are only superficial, in the sense that they can all be explained by a vast difference in degree of structural complexity between the human brain and the brains of other animals, then the materialistic dogma has obvious grounds in support of its prediction.

The present differences in the degree of structural complexity between the human brain and that of artificial intelligence machines can certainly be overcome in the future. There is no reason to suppose that machines cannot be constructed with parallel processing and with structural components and connections in excess of 10^{11}, which is the measure of the human brain's structural complexity. Artificial intelligence machines will then be as intellectually competent as we are. It may even be that the performance of these future machines will clearly surpass the best human efforts and accomplishments.

At present, the most powerful and intricate machines can do many things that brute animals cannot do, as well as many things that human beings can do, such as mathematical calculations, all forms of logical reasoning, and heuristic formulations. In addition, those machines can do these things better than many human beings, and more quickly than most. Their use in putting a man on the moon is a striking example of this. That could not have been done without them.

However, they cannot do some of the things that almost all human beings can do, especially those flights of fancy, those insights attained without antecedent logical thought, those innovative yet nonrational leaps of the mind.

Even the most optimistic computer technologists are willing to admit this while they remain confident in predicting the success of their efforts in the future. That confidence rests on their dogmatically asserted materialistic assumption that everything depends on the size and structural complexity of the brain and nervous system in human beings and of the material components and connections in the machines they hope to build in the future.

The declaration of the seventeenth century by the French philosopher René Descartes that *matter cannot think* is the battle cry of those who deny that the technologists' prediction will ever come true. The technologists' battle cry is just the opposite: matter can be made to think in all the ways that human beings think. Their thesis is the very opposite of the thesis of this book: not mind over matter, but matter over mind.

Since the position one takes on this issue depends on one's position regarding the intellect's immateriality, it may be used to open this chapter on artificial intelligence with answers to the objections that have been and can be raised against the thesis that the brain is only a necessary, but not the sufficient, condition of conceptual thought, and that an immaterial intellect as a component of the human mind is required in addition to the brain as a necessary condition.

First objection and reply. The clinical data of brain pathology, especially brain injuries that are accompanied by disorders of speech and by the loss of understanding, show the involvement of the brain in the processes of conceptual thought, just as other brain injuries causing blindness or deafness show the involvement of the brain in perceptual processes. Hence the one set of processes like the other must be a function of the brain. This objection was met by Aquinas in the thirteenth century. He dealt with the impediments to conceptual thought that result from brain injuries as well as the interference that results from the effect of toxic substances and fatigue poisons on the action of the brain.

Aquinas pointed out that there is no inconsistency between admitting the involvement of the brain in conceptual thought and asserting the intellect's immateriality. All that the evidence from brain pathology shows is that the brain is a necessary condition of conceptual thought, and in order to deny that the brain is the sufficient condition of conceptual thought, one does not have to deny that it is a necessary condition.

The error of the objection consists in treating conceptual and perceptual processes as wholly alike in being functions of the brain—that is, in treating *visual* blindness (loss of sight) as if it were the same as *conceptual* blindness or agnosia (loss of understanding). To treat them as the same is to ignore the argument for the immateriality of conceptual thought. The objection can hardly invalidate an argument that it ignores.

Second objection and reply. The human infant is not born able to exercise the power of propositional speech. It is only in the course of maturation that that power comes into operation and develops with exercise. The infant's first use of names or designators and his first utterance of sentences do not occur until, with growth, his brain reaches a certain magnitude. Hence it would appear that there is a critical threshold in the continuum of brain magnitudes above which the human being has and below which he lacks propositional speech. Since the presence of propositional speech is our only objective evidence of the presence of conceptual thought, it can be argued that engaging in conceptual thought depends, as engaging in propositional speech also depends, on a certain brain magnitude.

The reply to this objection, like the reply to the preceding one, concedes that conceptual thought depends on the brain, and especially on its having a certain magnitude. However, all this shows is that the brain, or a certain magnitude of it, is a necessary condition of conceptual thought. The argument for the immateriality of conceptual thought, the whole point of which is to show that the brain is not the sufficient condition of conceptual thought, remains untouched by this objection.

Third objection and reply. It has been conceded that animals and machines are capable of perceptual abstractions. Rats can learn to react to different triangles as *if* they all had some characteristic in common (their triangularity) that is not shared by other visible shapes; some success has been achieved in getting machines to recognize different shapes in an apparently discriminating manner (i.e., react in one way to square shapes, and in another to triangular shapes). It would thus appear that animals and machines are able to apprehend universals—classes or kinds of objects. But unless an immaterial power is to be attributed to subhuman animals and to machines, it would seem to follow that an immaterial power need not be posited to explain man's apprehension of classes or kinds of objects. Hence, even if it is granted that the concepts whereby we know kinds or classes are universal intentions, that does not justify our positing the immateriality of the power of conceptual thought.

The reply to this objection hinges on preserving the distinction between perceptual abstraction and concept formation. A perceptual abstraction, as attained by men or other animals, is an acquired disposition to perceive a number of sensible particulars as being of the same kind or as sufficiently similar to be reacted to in the same way; it is also a disposition to discriminate between similar and dissimilar particulars. It is not a disposition to recognize a single perceived particular as being of a certain kind, for the recognition of a single perceived particular as being of a certain kind is inseparable from the understanding of the kind itself. These related acts of recognition and understanding presuppose more than perceptual abstraction; they presuppose concept-formation. For a laboratory rat that has learned a food cue, a perceptual abstraction or generalization enables it to perceive that this shape and this shape (e.g., triangular shapes) but not this shape or that (e.g., circular shapes) are sufficiently alike to serve as the cue for a certain response.

But such perceptual generalization and discrimination does not dispose the rat to recognize that this shape by itself is a triangle or to understand triangularity when no triangular shapes are perceptually present. Only man, having the concept of triangularity, can recognize this perceived shape as being an instance of triangularity, and can, in the absence of any perceived shape, understand triangularity and the distinction between it and circularity.
By means of a perceptual abstraction, like that attained by the laboratory rat, man can also perceive a number of sensible particulars as similar shapes ' and discriminate between them and dissimilar shapes, but his recognition that the similar shapes are all triangles and that the dissimilar shapes are circles derives from his concepts of triangle and circle, which operate in conjunction with his perceptual abstractions.*

*Without concepts, we would only perceive, as animals do, the individual thing. If we reacted to a number of individually differing things in the same way, we

> would not be cognizing what is common to them or knowing them in their universal aspects; we would only be reacting to them as functionally equivalent stimuli. By means of concepts, and only by means of concepts, we understand kinds or classes as such entirely apart from perceived particulars and even though no particular instances exist. By means of percepts alone—*if that ever occurs in human cognition*—we would apprehend individual things without any understanding of them. This is the meaning of Kant's statement that percepts without concepts are blind, and concepts without percepts empty. Hence, if we are right in thinking that men have and other animals lack the power of conceptual thought, then we must also assert a difference in kind between perceptual processes in animals, which are blind in Kant's sense, and perceptual processes in men, which are enlightened by concepts.

The central point here is that perceptual abstractions do not function in the same way in man, on the one hand, and in nonlinguistic animals and machines, on the other hand. In man they operate in conjunction with concepts; in other animals and machines, they do not. It is only through concepts that we are able to understand kinds or classes of objects, and it is only through concepts in conjunction with perceptual abstractions that we are able to recognize this perceived object as being of a certain kind or class that we understand.

Perceptual abstractions by themselves, functioning in the absence of concepts as they do in animals and machines, can do no more than enable the animal or machine to discriminate between perceived particulars according to whether they are sufficiently alike or sufficiently different to warrant a particular reaction.

It might be difficult for anyone to say which of three technological innovations in the twentieth century, other than the harnessing of atomic energy, has had the most profound effect on human life: the motor car, the airplane, or the computer. But I have no hesitation in saying that, with regard to our understanding of man and his mind, the computer is not only foremost but stands alone. The centurisold controversies about the questions and issues raised in the preceding chapters have taken a new turn because of the promise or threat of what might be accomplished by machines in the future.

Computers of the latest generation and machines that are devised to be embodiments of artificial intelligence (hereafter referred to as AI machines) have done remarkable things, and not merely in the performance of mathematical operations but also in playing games like chess, in problem-solving, in perceptual performances, in the processes of learning, and in making decisions. With regard to these accomplishments, three points should be reiterated.

First, computers are able to do what no animal, even those nearest to man in the scale of intelligence, can do at all. If the question that we asked about man and brute—Do they differ in kind or in degree?—were asked about computers and animals other than man, there could be no doubt about the answer: clearly in kind. This explains why technologists concerned with constructing Al machines have no interest in simulating animal behavior, only human performances.

Second, when computers do what human beings can do, they do it with much greater speed and much greater accuracy, and often in magnitudes of complexity that exceed the reach of human beings. It would have been impossible, for example, to put a man on the moon without the use of computers.

Third, the difference in degree that exists at present with regard to the size and complexity of computers and human brains is still vast. The number of cells in the human brain is estimated to be 10^{11}, and the number of their interconnections is much greater than that. The number of transistor components in the largest computer so far constructed is 10^3 or 10^4 at the most. However, it is certainly both possible and even likely that computers built in the future will have componentry and interconnections in their circuitry that exceed what can be found in the human brain.

This last fact might lead us to suppose that when this happens, it should also be possible to build AI machines that will outperform human beings in all the things human beings can do and no other animal can do at all. If we jumped to that conclusion, we might then suppose that we had resolved the issue about whether the difference in kind between human and animal behavior was superficial or radical.

That supposition would be false because we had jumped too fast, we did not consider all the relevant alternatives. If the conclusion I argued for in the preceding chapter, as well as in the opening pages of the present chapter, is correct, then whatever progress is made in the future of computer technology, it will remain true that no Al machine ever to be constructed will be able to perform in a manner that is indistinguishable from human performance.

My argument, however, may have flaws in it, and I conceded that possibility when I said that while the contrary views cannot both be true, they can both be false. If my defense of the view that an immaterial factor must be present in order for there to be conceptual thought is not a sound one, then we cannot say it is impossible for an Al machine to be built that will be able to do everything the human mind can do.

Should that possibility ever be realized, the issue would be finally resolved beyond all reasonable doubt. We know that the Al machine is a purely material contraption. No immaterial factor enters into its construction. Hence, if it should demonstrate its ability to do everything the human mind can do, we would be compelled to conclude that the brain and nervous system with no immaterial factor added is not only necessary but also sufficient for all our mental activities, including the highest reaches of conceptual thought. That conclusion would carry with it the additional conclusion that the difference in kind between human beings and other animals is only superficial, not radical.

I hope I have made quite clear how crucially important is the role that Al machines can play with regard to the problems we have been considering. I hope it is also quite clear that, at this moment, we do not know which way the dice will fall. The failure of future Al machines to simulate every aspect of human behavior will be just as significant one way as their success will be the other way.

The future AI machine that will be put to the test of its power to *simulate* human behavior need not *replicate* the functioning of the human nervous system.* Unless "wet computers" are built, in which the impulses transmitted are electrochemical, not purely electrical, achieving the same end result but not achieving it in the same way will suffice.

*We must distinguish between simulation and what is called "replication." The attempts to construct mechanical models that operate in the same way the human brain operates are efforts at replication. In contrast to replication, the simulation of human behavior by machines consists in achieving the same end result in the way of performance but not achieving it in the same way. Thus, for example, airplanes simulate the flight of birds, but the mechanics of flight are not the same in both cases, though both bird and airplane obey the same laws of aerodynamics.

The fact that the AI machine is not a living organism, does not laugh or weep, is not subject to moods, nor manifests any nonintellectual aspects of human conduct has no bearing on the outcome of the test. We are concerned here only with how intelligent machines compare with the power of the human intellect, and not just with human intelligence.

The challenge to future technologists is very precise. It does not call for the production of an Al machine the performance of which will provide us with an answer to the loose and unclear question, "Can machines think?" In its use by psychologists, neurologists, computer technologists, and philosophers, the word "think" has so many meanings in its application to animals, man, and machines that if anyone asks "Can animals think?" or "Can machines think?" the answer will be "yes" in some senses of the word and "no" in others.

Fortunately, in order to make a critical test of artificial or machine intelligence, it *is only necessary to do what is possible*—namely, gain general acceptance of a definition of human thinking in all its variety. A test can be devised that involves a distinctively human performance, one that the Al machine must succeed in simulating. The human performance in question is that of conversation. As a defender of man's distinctive capacity for conceptual thought, I would be quite satisfied with that as a test of a machine's comparable intellectual capacity.

When Descartes declared centuries ago that matter cannot think, he challenged his materialist opponents to construct a machine that could engage in conversation with a human being. If that could be done, he was prepared to admit his error. I would say the same thing today. If computer technologists can succeed in constructing a machine able to engage in conversation with a human being, I, too, would admit error in the arguments I have so far advanced.

One of the most eminent mathematicians among the computer technologists of this century, Alan Turing, claimed that it is, in principle, mathematically possible to construct a machine able to pass this test. He went further and proposed a foolproof way to test his mathematically conceivable Turing machine.

Known as the Turing game, it involves substituting a machine for either a man or a woman in the game. When that game is played with a male and female human being, it is played in the following manner: a man and woman are placed behind a screen. An interrogator stands in front and asks them questions in order to see if he can detect whether the answers received come from the male or female behind the screen. The persons behind the screen are permitted to lie or to resort to any other trick that may help defeat the interrogator's effort. The answers are delivered in written form so that tone of voice does not reveal the gender of the responder. When this game is played properly, the interrogator's chance of succeeding is fifty-fifty, which in Turing's view indicates mere guessing and, therefore, failure for the interrogator.

In Turing's version of this game, an Al machine is substituted for one of the two human beings. Behind the screen is a human being and an Al machine. The interrogator asks them questions to try to differentiate the human from the machine responses, which are delivered, of course, in some uniformly printed manner. Those behind the screen must try to avoid detection by the way they answer the questions.

If the interrogator can succeed only at the level of chance (only 50 percent of the time), he is judged to have failed and the Turing machine is deemed successful. It has performed in a manner indistinguishable from that of a human.

A slightly altered and, perhaps, simpler version of the Turing game would involve a lengthy conversation between two individuals, hidden from one another by a screen, with both questions and answers delivered in printed form. On one occasion the two individuals would both be human beings. On another occasion one of the two would be a Turing machine and the other would be the same human being.

If that person could not tell which of the two conversations involved another human being and which involved a machine, Turing's claim would be verified—that a machine can be constructed to perform in a manner indistinguishable from what is accepted as a distinctively human performance involving conceptual thought.

To construct a machine able to play the Turing game or to engage in conversation with a human being, a number of obstacles must be surmounted. In the first place, the machine's performance cannot be one that is completely programmed. Everything programmed is predictable; even a certain range of random behavior, if programmed, is predictable. But the Turing game involves questions from the interrogator that are unpredictable by any programmer; so, too, are the turns and twists in a long conversation conducted by human beings.

The Turing machine, therefore, cannot have the kind of programming used in an ordinary computer. It may have what Turing calls "infant programming"—some fixed connections like the innate reflexes possessed by a human infant at birth. The Turing machine must also be able to operate through flexible and random connections. It must, in other words, be able to learn and be teachable; beyond that, for the purposes of playing the Turing or conversational game, it must be able to acquire the use of a natural language, such as English, which includes common nouns that are names with universal significance.

At birth, human beings are endowed with the ability to learn from experience, and especially the ability to learn any language whatsoever, as well as the ability to think about any subject whatsoever within the range of all possible thinkables. If such native endowment is regarded as programming, quite different from the kind of programming now put into computers, then a Turing machine that succeeded would have to be programmed in this way, not as computers are now programmed.

Holding the view that I have so far expressed—that the brain is not the organ of thought and that an immaterial factor in the human mind is required in order for its universal concepts to confer significance on the general terms we use in ordinary discourse—I regard it as highly improbable that a Turing machine will ever be built that is able to succeed in passing the Turing test.

Highly improbable does not mean impossible. To get at the truth about this matter, it is of the utmost importance for computer technologists to keep on trying to produce the machine that Turing claimed is, in principle, mathematically possible.*

*On this point, see the argument to the contrary in a book by a mathematical physicist who is also a philosopher: *Brain, Mind, and Computers* (1969) by Professor Stanley L. Jaki. The book was awarded the Le comte du Nouy Prize for 1970.

Philosophical arguments frequently fall to persuade the opponents at which they are directed. I have good reason to doubt that the arguments I have directed against the moderate materialists at the end of the preceding chapter and at the beginning of this one will persuade them that the brain may be the necessary, but cannot be the sufficient, condition of conceptual thought.

How else might they be persuaded of the soundness of my view, if in fact it is correct? By the failure of the computer technologists to build the requisite AI machine—one that can perform as Turing claimed it should be able to. Each time the technologists try and fail, the possibility of their success becomes less and less probable. Each successive failure increases the probability that machines will never be able to perform in a manner indistinguishable from that of the human mind.

The answer may not be forthcoming in the immediate future, but the pursuit of truth is an unending process involving the whole of available time. Up to the present moment, machines have not turned in the requisite performance. Reasons can be given for thinking that they never will. I am content to let matters stand that way. I hope my readers feel the same.

We welcome your comments, questions, or suggestions.

THE GREAT IDEAS ONLINE
is published weekly for its members by the
CENTER FOR THE STUDY OF THE GREAT IDEAS
Founded in 1990 by Mortimer J. Adler & Max Weismann
Elaine Weismann, Publisher and Editor
Phone: 312-943-1076
Mobile: 312-280-1011
David S. Peterson, Manager

A not-for-profit (501)(c)(3) educational organization.
Donations are tax deductible as the law allows.