[6]

DO COMPUTERS THINK?

. . . men say that the gods have a king, because they themselves are or were in ancient times under the rule of a king. For they imagine, not only the forms of the gods, but their ways of life to be like their own. ARISTOTLE, *Politics,* 1252b.

1. Introduction

THE SIMPLEST ANSWERS to the question whether computers think are, of course, the following: (a) No, machines do not and cannot think, because thinking is a spiritual activity, and spirit is altogether alien to matter; (b) yes, machines think, as shown by the fact that modern digital computers are able to perform the highest mental operations, which are the mathematical ones.[1]

The first answer is based on the dualistic tenet: to the upholders of this view, the claim that machines think is nothing short of a heresy. The second answer is usually not justified theoretically by those who utter it: they confine themselves to offer as a proof the obvious fact that computers perform complex operations which, *when performed by human beings,* are usually described as mental operations. Although no philosophical justification of this answer seems to have been advanced, it seems to be tacitly founded on the phenomenalist maxim that

[1] See, e.g., E. C. Berkeley: *Giant Brains* (New York, John Wiley & Sons, 1949).

things are exactly as they appear to be, and on the pragmatist view that what counts are net results.

Those who do not believe that things are ultimately simple cannot endorse such answers. The dualist reply blocks every advance in the direction of building machines designed to save physical and mental toil and which, in an adequate social setting, would contribute to attain what Wiener[2] has called a human use of human beings—and, let us hope, no *use* at all. As to the behaviorist answer, it may be and is actually being employed for devaluing creative work in favor of the routine of ready-made operations; besides, it does not foster the inquiry into the mental aspect of human life, as it usually takes for granted that no such mental aspect exists; finally, it brings us back to prescientific days, to the extent to which it assigns human attributes to inanimate objects.

I think that both solutions are dogmatic. The first, in so far as it assumes without evidence the tenet of the absolute heterogeneity of substances; the second, because it is an uncritical reply based on mere analogies. It seems necessary to try and look for a more satisfactory solution of this important problem, a solution based neither on the a priori rejection nor on the uncritical acceptance of phrases like "The electronic brain will think for you," "Computers may integrate differential equations," "The reading machine is able to abstract," and other items of advertising. What I propose to do here in order to ascertain whether machines think or not, is to examine succinctly the two main aspects of the question, namely, (a) the nature of computers, and (b) the nature of mathematical thought.

---

[2] N. Wiener: *The Human Use of Human Beings* (Boston, Houghton Mifflin Co., 1950).

2. Ideas and Their Physical Marks

Cyberneticians rightly claim that it is impossible to realise the scope of what they like to call the second industrial revolution while retaining the traditional idea that machines are completely passive tools in the hands of craftsmen. They are right in maintaining that only an examination of artifacts endowed with a high degree of automatism can lead us to a reasonable evaluation of cybernetics. But it seems that they usually neglect the second side of the problem, viz. that of the nature of thought processes and objects, in particular of those which computers are said to handle. And, since machines are designed to mimic thought, a misunderstanding of the latter's nature will produce, by a sort of (positive) feedback, a misunderstanding of the nature of the very machines designed by cyberneticians to replace some mental functions.

It is certainly true that, in so far as machines are the outcome of intelligent and purposive work, they cannot be put in the same class as inanimate objects; machines are *matter intelligently organised* by technology, and as such they stand on a level of their own. But, on the other hand, it should be kept in mind that artifacts, however complex, operate only with material objects, never with ideal, abstract objects—a sort of operation which is precisely one of the distinctive characteristics of educated human beings. This elementary point is missed by most cyberneticians, and it seems to be the clue for the understanding of the whole question. Indeed, the confusion between "thinking machines" and "machines that replace thinking" lies in the identification of mathematical objects with their materialisations and, in general, in the identification of concepts and judgments with physical marks representing them. Once such an identity between sign and designatum has been accepted, once such a confusion of the material and the ideal levels has been indulged in, it goes without saying that machines think.

We can express and record thoughts by material means, so that material objects (pennants, knots, acoustic or electric signals, spoken or written symbols, etc.) are correlated with them. Such physical marks correspond to thought, they stand for thoughts, they represent thoughts, they are deputies of thought—but they *are* not thoughts. If, instead of relatively static materialisations of concepts and ideas, we employ devices that combine and transform such physical marks (whether they do it automatically or not is immaterial to our concern), then we will have constructed artifacts that do not think but which, up to a certain point, can represent and thus replace human thought within certain limits. All this is elementary, but seems to have been forgotten by most apologists of cybernetics, who systematically confuse ideal objects with their physical correlates.

I wish to emphasise the following points: (a) the physical processes organised by technology and involving material representatives of ideal objects, are *correlated* or coordinated with reasoning, although they are not rational: what is at stake is not an identity in kind, but a similarity of pattern; (b) those physical processes depend on the nature of the machine rather than on the nature of thought, as shown by the fact that a given mental process can be materialised in several different ways; that is to say, up to a certain point the physical marks used to represent ideas are contingent upon the latter's nature, and they do not depend on the context in which such ideas appear, since, when designing machines, technologists are interested only in external similarity, in resemblance of structure; (c) such physical processes represent combinations of ready-made and clear-cut thoughts, but they are not able to create representatives of thoughts that do not arise as mere combinations of old thoughts in accordance with the rules of logic built in the machines.

All this is clearly illustrated by pencil and paper operations, for the problem of automatism, while central for technology, is rather irrelevant to the question whether machines think or not. Indeed, we are not asking whether machines think on immediate or on long-run command, but whether they think at all. We might as well take a step beyond, considering the first and simplest computer, namely, the abacus. But, since I do not wish to humiliate technicians, let me recall an artifact of a more evolved type, viz., Pascal's *machine arithmétique* (1643), which is a sophisticated abacus.[3] In this artifact each integer, from 0 to 9, is represented by (or materialised in) a cog of a cogwheel; there is one cogwheel for units, another for tens, and so on. When the first wheel turns one-tenth of a complete revolution, it has "counted" 1; the gear wheels are connected in such a way that after 10 such unit rotations—i.e., after the machine has "counted" up to 10—the second wheel turns automatically one-tenth, thus recording 10 units in one stroke, while the first wheel regains the zero position—and so on. In this mechanism, which is substantially the same still used in desk-computers, the mathematical operation of addition is represented by the physical process of rotation of wheels through definite angles.

[3] B. Pascal: "Advertisement necessary to those who have curiosity to see the Arithmetic Machine, and to operate it," in D. E. Smith (Ed.): *A Source Book in Mathematics* (New York and London, McGraw-Hill Book Co., 1929).

Computers differ according to the sort of material recording employed in them. But, irrespective of their degree of automatism, they are all characterised by the fact that *they do not perform mathematical operations,* but only physical operations which we coordinate with mathematical ones. This holds not only for the computers of the analogue type (such as the slide rule), but also for digital computers—in spite of the fact that it is usually stated that in the latter numbers are operated upon directly. They all perform physical operations on entities (cogs, electric pulses, switches, etc.) that record ideal entities, that represent them at the level of technology; in this respect, the essential difference between analogue and digital computers is that the latter, in contradistinction to the former, operate with denumerable (discrete) events—but not with numbers! Herein lies the decisive difference between natural inanimate objects and artifacts: the former are not the materialisation of images, concepts, ideas, etc., whereas artifacts *and other* concrete culture objects—such as books, paintings, phonograph records, etc.—do represent facts of mind in a material form.

3. Counting

The most advanced computers are at present those of the digital type, that is, those based on coordinations of sequences of discrete material events (such as electric pulses). They perform operations which, when performed by man, are called arithmetical—i.e., addition, subtraction, multiplication, and division, as well as recognition of sign and equality of numbers. They also perform physical correlates of operations which, while not being arithmetical in kind, can be numerically approximated by sequences of arithmetical operations; for example, integration of most functions met with in practice may be numerically approximated, to any desired degree of accuracy, by summation.

By means of adequate devices, the four elementary operations can be reduced to addition; and addition of natural numbers and of ratios of them (which are the only ones recorded in digital machines) is after all reducible to counting, since any integer is obtained out of the preceding one by means of the operation $+1$, which is the typical operation performed in counting. Hence, in the last analysis, digital computers, working as they do with physical correlates of integers, are *counters at the physical level.* The basic traits of this kind of counting are: (a) it works on the basis of a coordination of *physical* entities; (b) computers are *specific,* that is, they do not count anything that is countable, but only very particular objects (angles, electric pulses, etc.).

Counting at the physical level is certainly analogous to the way primitive, or modern uneducated man, counts; indeed the latter use a correspondence between the things they want to count, and their fingers; i.e. they coordinate two sets of material elements. But, whereas in the machine the coordination remains at the physical level, in man the connection between the two sets (e.g., shells and fingers) passes through the central station of intelligence, and this is what enables man to count whatever he chooses to (provided it is not continuous).

But there is a higher level of counting. When mathematicians count, they ususally operate neither like computers nor like uneducated men. What the mathematician means by counting is the process characterised by the following features: (a) to count is to establish a correspondence between the given set of objects and the set of *ideal* objects called natural numbers; this is what enables us (b) to count, at least in principle, every countable (discrete) set, whatever its nature may be; that is, mental counting is not specific, like physical counting, but *generic*; (c) besides, the mathematician knows not only how to count but also *what* counting is; for example, he is able to define counting as the coordination (one-to-one correspondence) between the given set, and the set of natural numbers; finally, (d) he knows *why* he is counting.

There is little room to doubt that machines do not attain this latter level of the operation of counting: they neither use abstract objects, nor are able to count everything, nor know why they work. Of course, for so-called practically-minded people such things are irrelevant, the important point being that machines count in some way. But this pragmatic attitude, besides being unphilosophical, may finish by blocking further technical advances in the building of servomechanisms, for it clearly entails the assumption that technology should tend to mimic mind at its lowest possible levels—which is a good norm of technological strategy during the first stage, but is obviously shortsighted.

4. Adding

Machines add concrete numbers; they do not add just numbers, i.e., pure, or abstract numbers; hence machines do not add in the usual sense of the word. In nature addition takes place in an infinity of concrete manners, as contrasted to arithmetics, where addition of numbers is univocally defined. Unlike nature, and unlike artifacts, man is able to build a mental level upon the physiological one; his cortex is able to perform what no assemblage of valves, relays, and switches can do, viz. the feat of adding abstract, or pure, numbers, which need not refer to things outside them.

There does not seem to exist any doubt that our brains perform abstract operations by means of very concrete physiological processes, just as physiological functions have in turn an underlying physico-chemical basis. Moreover, since in most cases our capacity to operate the abstract machinery is limited, we are forced to take pencil and paper, or other assistants—such as computers—in order to manipulate mental entities in a material way—a method which facilitates automatism and makes a larger use of the senses. That is, we handle ideal objects by means of material operations taking place within our skull, with or without the additional help of external materialisations of those ideal objects. In this way we are able, at least in principle, to add all imaginable things: not only angles but also angels, not only electric impulses but also emotional impulses. We are always able to translate concretely found numbers into abstract numbers, and vice versa. The "nervous computing machine," as Wiener calls our central nervous system, works not only at the physiological level but also at the logical level, which has laws of its own. This truism seems to have been forgotten by cyberneticians, who hold that computers operate with numbers,[4] i.e., with ideal objects, whereas on the other hand they claim that the operations of the mind are reducible to electric terms, so that brains would "ultimately" work only at the physical level.[5]

---

[4] N. Wiener: *Cybernetics* (New York: John Wiley & Sons; Paris: Hermann et Cie., 1948), p. 136.

[5] W. Grey Walter: "An Electro-Mechanical Animal,'" in M. Monnier (Ed.), *L'organisation des fonctions psychiques* (Neuchatel, Ed. du Griffon, 1951). W. R. Ashby, *Design for a Brain* (London, 1952).

Machines, let us repeat, do not add pure numbers; they add turns of cogwheels, electric pulses, etc., which combine in accordance with physical laws peculiar to each such material objects. There is, to be sure, a parallel between some of these physical laws (e.g., those of the current flow in circuits) and certain logical laws (e.g., the laws according to which propositions are combined); computers are based on this isomorphism, but this does not involve an identity in kind. Machines are not, as our nervous system is, multilevel structures; they cannot retranslate concrete objects into their abstract correlates or vice versa. It is we who perform such a translation, when building and using the machine. We do it whenever we insert the input "message" (i.e., physical correlate of information plus operational symbols) and when we collect the output "report."

In other words, the operator has to perform at least the following operations: (a) to translate a group of abstract (mathematical and/or logical) entities into the physical "language" of the computer; (b) to retranslate the output "message" into the abstract language of mathematics and/or logic. What the machine does in our place is the intermediary stage of "information processing." The amount and quality of mental work required to handle computers (not to speak of their design) is such, that a specialist has written that, "Perhaps, if IBM's familiar

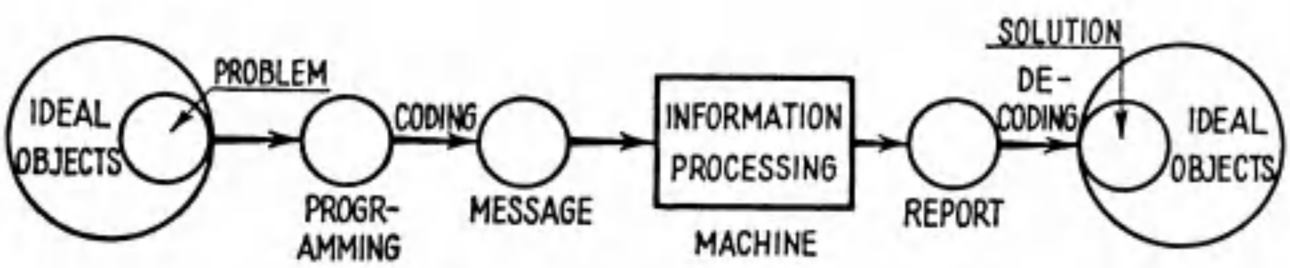

Figure 3. The computer: an auxiliary link in a long chain.

motto ['Think'] needs amending, it should be 'Think: Think harder when you use the *Ultimac.*' "[6]

If one avoids using figurative language—and, particularly, if one avoids Wiener's mistake of assigning a language to machines—one is forced to admit that, strictly speaking, machines are neither fed with "information" nor "work with logic," nor "report" the results of their work, but that they are "fed" with physical marks (e.g., taping) and yield another group of physical marks of the same nature (output) which, when interpreted or read by the operator, get converted, in his brain, into authentic information.

The foregoing remarks apply to all of the operations performed by the computers of the digital type, since such operations are all reducible to counting and adding; and most of our remarks apply also to other types of automata, for they all execute purely physical operations after a certain programme.

5. Pythagorean Machines

This is perhaps the place to point out an important limitation of digital computers, which cyberneticians, enthusiastic as they are over numerical results, tend to overlook: it is the fact that digital computers are utilisable only if the given problem is numerically "reducible" to a succession of arithmetical operations involving solely rational numbers (integers and fractions). For example, they do not integrate, but add; they do not yield square roots, but fractions approximating them. The differences between the obtained and the exact result may be negligible from the *quantitative* point of view, but

[6] A. L. Samuel: *Proceedings of the Institute of Radio Engineers, 41*:1223 (1953).

they are enormous from a *qualitative* point of view. For instance, in the case of integration, an entity—infinity—and an operation—the approaching to the limit—are lost.

This shows the enormous limitation of digital machines: they are, to speak the pictorial language that cyberneticians love, Pythagorean machines, for they are limited to counting—at the physical level. They, not mathematics, fit into Mach's definition of the latter as "the economy of counting."[7] Since computers work with discrete strings of events, they "ignore" the continuum; and, working with actuals, they remain on this side of infinity. But irrational numbers and infinity are just some of the most characteristic objects of classical mathematics, as contrasted with ancient mathematics. If mathematics is the science of infinity, as Weyl has claimed, then it is plain that computers, imprisoned as they are in the narrow frame of material representatives of natural finite numbers, do not perform mathematical work.

Practically-minded people tend to conceive of mathematics as The Art of Computing. No wonder, then, that they should believe that computers perform mathematical (and logical) work. Now, even granting that computers calculate (which, as we have seen, is not true), the truth is that computation does not exhaust mathematics, just as the logical calculi do not exhaust logic. Computation is, to speak loosely, the "mechanical" side of mathematical work; computation is a part of mathematics which is concerned neither with creating mathematical objects, nor with framing the rules in accordance to which they are operated upon, nor, of course, with their metamathematical examination. A computer,

[7] E. Mach: *The Science of Mechanics,* transl. by T. J. McCormack (La Salle, Ill., and London, The Open Court Publ. Co., 1942), pp. 583, 584.

whether human or artificial, need not know the nature of the numbers with which he or it operates, nor to worry about the meaning and foundation of the rules governing the combinations of numbers: both numbers and rules are given to him or it, and he or it proceeds to use such a material in a blind, "mechanical," way, without worrying whether natural numbers are primitive concepts or classes of classes. Computers, in short, do not need to "know" what numbers are, but only how to work with their concrete correlates (e.g., holes on a punched card).

This does not mean that machines are equivalent to human computers; although some net results of computing machines may be correlated with some net results obtained by human computers, the difference between them is as big as between a striped sweater and a zebra. For, whereas human computers know at least how numbers work—or at the very least how their written tokens work—computing machines do not *know* it—nor know anything else. In the first place, because they do not work with mathematical objects but with physical representatives of them. Secondly, because computing machines just perform certain operations without being aware of it: they do not know what they are doing nor even *that* they are doing anything—and this, simply because machines have no consciousness, which is a prerequisite for the non-automatic type of knowledge.

6. Are Machines Aware?

I have just employed a word which is the *bête noire* of behaviorists and cyberneticians, namely, 'consciousness.' No argument employing this term will consequently be accepted by them. However, it is difficult to see how certain problems can be solved without the help of the concept of consciousness; one of them is the problem of the various levels of knowledge. At least two levels of human knowledge must be distinguished: subliminal and conscious knowledge. Most of our knowledge, and especially most of the knowledge on which our everyday decisions and actions rely is of the latent or subliminal type, even in the field of scientific work; intellectual knowledge, on the other hand, is not intuitive but symbolic, and it is not automatic but conscious. Of course, conscious knowledge may finish by being automatised, and subliminal knowledge may emerge on the level of consciousness, but this does not efface their differences.

This is not the place to take the defence of knowing as the highest function of consciousness, nor of consciousness as the highest function of the person; the reader will therefore be asked to accept the fact of consciousness as a datum in the philosophical context—though as a central problem of human psychology—and to receive, just to prevent misunderstandings, an admittedly coarse definition of 'intellectual knowledge,' or second-level knowledge. It runs as follows: "Intellectual knowledge of X is awareness of the fact that X has such and such characteristics."

On this definition it is obviously wrong to hold that machines know anything in the intellectual sphere, for they lack consciousness. To say that a machine *knows* how to solve the problem stated in the programme is like saying that the planets *know* the Keplerian laws which they approximately follow, or that plants *know* how to photosynthetise. On the other hand, our definition ensures that mathematics belongs to the field of knowledge (though not of the external world), since it is practised by men aware of their own work—at least in the interesting cases.

If mathematics had nothing to do with concepts and were nothing but the blind performance of non-conceptual operations, as Hegel[8] thought, then it would not be a branch of knowledge. The Hegelian Croce[9] would then be justified in asserting that mathematics—which he called *simia philosophiae,* the ape of philosophy—has no cognitive value, but only a practical value. And the cybernetician McCulloch[10] would also be justified in writing that "We are to conceive of the knower as a computing machine."

Fortunately this is not so; mathematics belong to the highest achievement of human knowledge, and in order to be a mathematician—i.e., a knower of and a worker in some branch of mathematics—it is not only required to know how to calculate something (in the large sense of the word 'calculation,' i.e., meant as a combination of symbols), but also to know what and why it is being done. That is, the theoretical value of mathematics is great because it is not reducible to a set of calculi.

A mere glance at actual mathematical research—not, however, at certain books on the philosophy of mathematics—will convince anyone that this kind of work takes place at various levels, no one of them being entirely reduced to the combinatory level, i.e., to calculation—which is precisely what computing and "logical" machines mimic. Besides inference of the analytic type, mathematics contains two further layers: (a) a *synthetic* level, consisting of the framing of postulates, definitions, rules of operation, etc.—in short, all the concepts and statements be-

[8] Hegel: *Science of Logic,* transl. by W. H. Johnston and L. G. Struthers (London, G. Allen & Unwin Ltd., 1929), II, p. 322.

[9] B. Croce: *Logica come scienza del concetto puro* (Bari, Laterza, 1928), pp. 233-4 and *passim.*

[10] W. S. McCulloch: "Through the Den of the Metaphysician," *British Journal for the Philosophy of Science, 5:* 18 (1954).

longing to the level of principles; (b) a *critical* level, characterised by the critical examination of principles and theorems, and the clarification of their meaning–which may lead to the reconstruction of theories, i.e., to new stages of the synthetic level.

To use pictorial language, one may say that machines "apply" principles. But, so far as I know, nobody has as yet proposed to design a machine capable of building a new branch of mathematics, nor to make a criticism of whatever inconsistencies may be found in available mathematical theories. And this, in spite of the fact that cyberneticians claim that the processes of criticism and consequent correction are always "essentially" of the feedback type.

7. Can Induction Be Mechanised?

Available computers perform physical operations that are coordinated with logical and/or mathematical operations of the *analytical* type; they do not perform anything that can be correlated with analogy or with induction (whether ordinary or statistical). But it has been claimed that plausible reasoning, and in particular induction, can in principle be "mechanised": that it should be possible to build a machine capable of accepting or rejecting an hypothesis by examining instances of its consequences.

The possibility of "mechanising" probable (non demonstrative) inference seems more than doubtful, because the weight or strength we assign to the conclusion of plausible reasoning depends not only on stated reasons (assumptions and data) and on rules of probable inference that have been found successful in the past (and adopted on inductive grounds!); that weight depends also on a vague background of *unstated* "reasons." This background is constituted by our personal experiences, training, beliefs, and even hopes. That is to say, probable inference is not only based on clear and definite assumptions and bits of specific information, but also on a personal background (the mass of our *Erlebnisse*) and on a social background (our own *Zeitgeist*). And none of these backgrounds can be "fed" into a machine.

Polya[11] has expressed this thought with his admirable clearness: demonstrative and plausible reasoning have different tasks and appear from the outset as essentially different procedures: "demonstrative reasoning as definite, final, 'machine-like'; and plausible reasoning as vague, provisional, specifically 'human.' [. . .] In opposition to demonstrative inference, plausible inference leaves i n d e t e r m i n a t e a highly relevant point: the 'strength' or the 'weight' of the conclusion. This weight may depend not only on clarified grounds such as those expressed in the premises, but also on unclarified unexpressed grounds somewhere in the background of the person who draws the conclusion. A person has a background, a machine has not. Indeed, you can build a machine to draw demonstrative conclusions for you, but I think you can never build a machine that will draw plausible inferences."

Now, both in ordinary life and in science conclusive inference is only one side of thought. Powerful generalisations, such as those of factual science, are not framed along purely analytic ways; and fruitful analogies, such as the one between brains and computers, are not reached at through deductive chains. Moreover, presumably such in analogy could emerge and be worked out only in a

[11] G. Polya: *Mathematics and Plausible Reasoning* (Princeton, Princeton University Press, 1954), II, pp. 115-116. See also I, p. 198.

specific socio-cultural setting: namely, in our industrial civilization. Certain strains of thought can certainly be "mechanised" in the above-mentioned sense. What cannot be "mechanised" are inferential processes leading to proposals like "Thought can be mechanised." This requires the *souplesse* of informal logic, the persistence of mechanistic world views, and the audacity of contemporary technology.

In short, the *esprit géométrique* can be "mechanised" up to a certain point, whereas the *esprit de finesse* is basically unmechanisable.

8. Do Machines Abstract?

At this point the orthodox cybernetician might step in arguing that, as there are levels of mathematical work, there are also stages in the development of machine-building, so that one cannot be too sure that future artifacts will not surpass those of the analytic type. A reply to this objection could be: (a) No machine can ever attain the level of abstraction because machines "merely" *represent* abstract thought; they do not handle abstract entities nor, *a fortiori,* can they create new abstract objects, as they are secluded in the circle of inanimate matter, on which man can stamp his intelligence, but which lacks the material prerequisite to attain intelligence, namely life; (b) is it not much easier and important to beget and to train normal mathematicians?

Our hypothetical cybernetician would probably rejoin that, while it is true that computers so far built lack the capacity for abstraction, other machines have it. For example, the "reading machine" designed by McCulloch and Pitts is said to have such a faculty: it is able to "recognise" the same general shape, or pattern, in material objects having individual differences (e.g., printing-faces of different sizes and styles). Cyberneticians hold, in sum, that this machine "recognises universals."

The assignment of the faculty of abstraction depends, of course, on the meaning attributed to the word 'abstraction.' Also, it is plain that cyberneticians use in this connection the common, non-technical acceptation of that word, namely the one according to which abstraction consists of taking away, dispensing with, taking aside. Now, it should be remembered that this is not the sole connotation of the word in question. Moreover, to employ the word 'abstraction' to designate such an operation is often misleading, as it applies not only to mental but also to physical processes. Indeed, on that definition of 'abstraction' it might be said that the gravitational field has the faculty of abstracting in the highest degree, for it pulls all sorts of bodies, "abstracting from," or dispensing with, their properties. Would it not be nonsense to hold this? My claim is that the same kind of "abstraction" works in the so-called "recognitive artifacts" —not, however, the faculty of performing the abstract operations of the synthetic and the critical types referred to above.

In fact, what is the mechanism by which "recognitive artifacts" are said to "perceive abstract forms"—whatever psychologists may think of the possibility of *perceiving* universals? Essentially it is the principle of specific sensitivity (or insensitivity) operating, e.g., in wave filters, which "recognise" whole groups or bands of frequencies. This is not too distant from the humble sieve used in the kitchen to separate bodies of different sizes regardless of their specific nature. In all these cases only physical laws are involved, and not mysterious ones.

The claim that the "reading machine" (i.e., the artifact that converts optical into acoustical signals) is able to abstract, might be justified on the empiricist doctrine of abstraction. According to traditional empiricism, abstraction is only taking away, setting aside, ignoring, or forgetting particulars—never adding anything new; for this school, abstraction is mere schematic representation in thought of facts of experience. This doctrine—shared by detractors of intelligence, like Bergson—may account for the *first* level of abstraction, the one characterised by generalisation through elimination of particulars. This is the kind of induction that dogs perform when they learn to put all cats in a single enemy class; it is also the one we make when speaking of the cardinal number of a collection irrespectively of the nature of the elements of it. To this lowest level of abstraction, which we share with the higher animals, the usual definition of abstraction does apply. But this is not the sole level of abstraction attained by man; and the higher levels are not entirely reducible to the lower, although they are rooted to them.

9. Can Machines Outdo Their Designers?

Man is not only able to ignore or to forget—a privilege which he is not willing to share with machines, at least as regards the sphere of intelligibles. Man is not only able to disregard individual characteristics concentrating on common traits; he is also able to invent new objects not suggested to him, at least directly, by experience. For example, when we speak of moving bodies in general, we stand on the first level of abstraction; but when we refer to bodies and to motion separately, we perform a sort of quartering of sensibles, thus stepping on a higher level of abstraction. Again, when we introduce the concepts of actual infinity, irrational number, abstract space, energy density, vector potential, and the like, we create ideal objects lacking an empirical counterpart, although they may be correlated with experimental data via certain intermediaries: here we are moving on a third level of abstraction, the level of ideal objects not originated in simplification (first level) nor in quartering (second level). This third level of abstraction is characterised by new "emergent" qualities—although the followers of the empiricist tradition maintain that even "our most remote abstractions are all ultimately reducible to primitive atomic propositions and the calculus of the lowest level"[12] and that, in its turn, atomic propositions are nothing but peculiar nerve impulses.

Machines are not entitled to be even compared with their designers in the field of the higher levels of abstraction; as has been suggested above, some of them can "recognise" universals of the first degree (e.g., *squareness*) —in the same sense as a home-made hygrometer, lacking a graduated scale, might be said to "recognise" the universal *humidity*. The physical processes occurring in "recognitive artifacts"—and also in non-recognitive ones —are the material correlates of first level abstraction. The same as in the case of computers, what is at stake is a material representation of a mental function, not the function itself.

Obviously, machines are usually built because they can do some things which man either can do but painfully (washing machines), or slowly (differential analys-

[12] W. S. McCulloch: "Why is the Mind in the Head?," in M. Monnier (Ed.) : *L'organisation des fonctions psychiques* (Neuchatel, Ed. du Griffon, 1951) , p. 38.

ers), or inaccurately (lathes), or which he cannot do at all (airplanes, radio sets, piles). In this sense machines surpass their builders, thereby falsifying the scholastic maxim (adopted but not invented by Descartes) that there can be nothing in the effect that had not been in some way in the cause. But machines cannot surpass man in everything, even though we are told that the new computers "are capable of learning and thinking far beyond us."[13] For all their usefulness, machines are products of culture whereas their designers and builders are, besides, producers of culture objects. And, rigorously speaking, machines surpass nobody in nothing: what happens is that a given designer of machines may surpass himself, or some colleague of his by building an improved machine.

Perfectibility is indeed a characteristic of living matter absent in machines. Perhaps some machines can "learn" something, i.e., can be said to have an experience and to make use of it. But, (a) theirs is, so to speak, an individual perfectibility, since it is not transmitted to the species *machina ratiocinatrix* through reproduction or through culture: it begins and ends in the individual machine; (b) machines do not seem to be able to advance in a sense very different from the way animals progress, namely by trial and error; this is, indeed, the behavior typical of machines with self-correcting (feedback) mechanisms, and is the least intelligent way of learning, because it is not planned and because it does not make use of another's experience; (c) rigorously speaking, machines do not "learn" by themselves, but are

[13] W. S. McCulloch: reference 10.

"taught," either by their designer or by external circumstances.[14]

Man learns not only as an animal, i.e. through individual mistakes: he learns mainly through the agency of society, which acts on his biological and psychological mechanisms. This is why man can dispense, to a large extent, with purely biological progress, advancing at a rate that is without a parallel among lower animals. One of the reasons why man covers levels of learning higher than the peculiarly animal level, is that he is endowed with consciousness: unlike the machine, man is able to know what he does, how he does it, and why he is doing it; he is, moreover, able to foreknow his doings. Therefore man comes to know, among other things, that he must go forward in order to survive; and in some cases he is even able to discover that he himself deserves the credit for it.

10. Artificial Thought?

All machines save both mental and physical toil. But they do not always save work because they do it in our place. For example, a (new) car may save us the effort of walking, but not because it walks instead of us; the car performs a completely different operation, which amounts to walking only in so far as both motions have the *net result* of displacing our bodies over space. The same holds for computers and other "machines that think": to assert that they think is as erroneous as saying that cars walk. Machines do not save us mental work because they do it, but in spite of the fact that they do

---

[14] M. V. Wilkes: "Can Machines Think?," *Proceedings of the Institute of Radio Engineers, 41*:1230 (1953).

something very different, which the designer has *correlated* with certain mental operations.

In this very restricted sense, computing machines may be said to perform what has been called *artificial thinking*.[15] Not in the same sense as synthetic compounds, such as vitamins, are called artificial, for the properties of the artificial and of the natural chemicals are often exactly the same, which is obviously not the case with artificial thought. (Think of the punched card yielded by some computers.) Machines can be said to perform artificial thinking in the same sense as cars can be said to perform artificial walking: because they yield net results which are equivalent to the model in a single respect—whereas in the case of synthetic compounds the identity often covers all known aspects.

Mays[16] has coined an irreplaceable formula for designating "machines that think": he said that they *think by proxy*. The full meaning of this statement should be appreciated, especially since it is metaphorical. To say that digital computers think by proxy does not mean that they think only in a limited way, or lazily, or solely on command—not even that they think for us, nor for our sake. It means that they do not think at all, although they perform operations that represent our thought in a certain field, yielding results devoid of intellectual content but which, when translated into the language of ideas, can usefully be incorporated in reasoning. To marry by proxy may have a legal value, but no more than this; something similar happens in connection with machines: man does not delegate thought to the computer,

[15] P. de Latil: *La pensée artificielle* (Paris, Gallimard, 1953).

[16] W. Mays: "The Hypothesis of Cybernetics," *British Journal for the Philosophy of Science, 2*:249 (1951).

for the simple reason that the computer cannot think, but can instead perform functions which we correlate with thought. Analogously, a portrait may represent a person, but it is not a person; to confuse both may lead to magic.

The computer, like every other automatic machine, is made to run for our sake; it would be wrong to infer from this that it acts *like* we. To commit this fallacy—and most cyberneticians indulge in it—is the same as to confuse the deputy with the deputised thing. This is what people do when they confuse a piano-player with a pianist, or the vicars of God with God. This fallacy, of inferring that something acting for us must in some way or other participate in human nature, is typical of primitive and archaic logic: it is called reasoning by participation, and is the kernel of magic rituals. To conclude essential kinship in nature from mere correlation, from resemblance in pattern, is to push analogies too far; so far, that the difference between science and magic is lost. Needless to say, science began when that very procedure of too many cyberneticians, namely the magical play with anthropomorphic analogies and with metaphors, was rejected.

11. Metaphors and Their Misuse

A distinctive mark of cyberneticians is their love of metaphors. Thus, they use to say that artifacts think, know, receive and supply information, learn, wish, and even get sick. This is one of the main troubles with cyberneticians, namely, that they usually fail to distinguish between identity and resemblance, between the model and the portrait; that, in short, they use key concepts in wrong contexts. When a whole science and a whole philosophical literature are built on linguistic traps, one is entitled to distrust the slogans by means of which the new faith is advertised—or, at least, one has the right of demanding a purification of language, and the right of smiling at certain warnings[17] against that very use of concepts out of their proper context, in which so many cyberneticians systematically indulge.

However, if the confusion of somebody with his deputy may lead to nonsense, it would be equally foolish not to realise that sometimes there may be something in analogies. Two very different objects may have something in common at some level or in some respect—and usually material objects do have a lot of features in common. To realise this is as important as to avoid concluding sharing of essentials from mere resemblance in particulars or even from similarity in structure. For example, memory in computers and in man are assuredly totally different at the physiological and at the psychological levels, at which machines do not even exist; but there is an analogy (similarity of pattern) at the physical level, for what is properly called 'memory' in the case of higher animals, and improperly so in the case of artifacts (where it might be called 'storing'), is the capacity of retaining or storing some condition (whether in a static or in a dynamic way). Not to recognise such general traits shared in common, or likenesses in pattern, may lead us to support dualism or idealism with regard to the so-called mind-body problem, thus favoring the return of the much discredited philosophical (or literary) psychology, still in vogue in Germany and its philosophical de-

[17] N. Wiener: "Some Maxims for Biologists and Psychologists," in M. Monnier (Ed.): *L'organisation des fonctions psychiques* (Neuchatel, Ed. du Griffon, 1951).

pendencies. But to claim that partial identities and formal resemblances are *all* that matters—holding, for example, that machines can store ideas—is to push analogies so far that their heuristic function becomes lost sight of.

Now the whole of cybernetic literature is infested with such physiological and psychological analogies. The fact that some of them are deep and fertile lends it strength; the fact that they are nothing but analogies deprives cybernetics of methodological solidity—at least in the opinion of those who do not accept the philosophy of the *as if*. The great merit of cybernetics lies, in my opinion, in having pointed out and worked out something which was far from new but which is true, namely, the physical *basis* of life and mind functions. The main shortcomings of cybernetics are probably, (a) to have proclaimed that life and mind have no such physical basis, for they *are* just physical phenomena (mechanistic levelling down), and (b) to have levelled computers up to the level of the human nervous system (animistic reduction).

The levelling down is effected by way of what has rightly been regarded[18] as the central hypothesis of cybernetics; according to it, the essential mechanism of the nervous system is a purely physical one, namely negative feed-back. The levelling up lies in the claim that there is no distinction in principle between the observable behavior of a suitably designed artifact, and the behavior of the human brain.[19] This peculiar blend of animism and mechanism, which characterises the cybernetic literature,

[18] J. O. Wisdom: "The Hypothesis of Cybernetics," *British Journal for the Philosophy of Science*, 2:1 (1951).

[19] D. M. MacKay: "Mindlike Behaviour in Artefacts," *British Journal for the Philosophy of Science*, 2:105 (1951).

might be called *animechanism* or, as has recently been proposed, *technozoism*.[20]

To say it in fewer words, the positive contribution of cybernetics consists, in my opinion, in its emphasis on the existence of connections between levels the very existence of which it denies—namely the physical, the biological, the psychological, the intellectual, and the cultural levels.

12. Conclusions

To sum up, we may say that computers count, add, etc., at the physical level, performing operations that are usually not regarded as mathematical (at least by mathematicians), since mathematics, an abstract science, is not interested in cogwheels, switches, electron tubes, electric pulses, etc. It is we who frame a correspondence (when building, "feeding" and reading the machine) between the concrete objects handled by the computer and our abstract objects. Without the human initial and final work of translating abstract into concrete objects back and forth, i.e. without the work of coding and decoding, the best of computers is helpless. In this respect, highly automatic machines do not differ essentially from the modest pencil, the simple abacus, or the cheap desk-computer, even though they are essentially different from a technological point of view.

Strictly speaking, computers do not compute, machines do not think, but they perform certain physical operations that we coordinate with certain mental processes. Since co-ordination, or one-to-one correspondence,

[20] H. Rodríguez: "Cibernética y pensamiento humano," *Congresso Internacional de Filosofía* (São Paulo, Brazil), III, p. 889 (1956). *Psicología y cibernética* (Buenos Aires: Siglo Veinte, 1958).

defines identity in pattern, the whole resemblance between machines and man is an identity of pattern, a formal identity or isomorphism of some of the operations of the machine and a small section of human activity. Without the intervention of man's abstract and purposive activity, which has no counterpart in machines, the most expensive digital computer is mere scrap iron.

Machines, however automatic, are tools, that is, material assistants of man. To hold that they compute, think, know, learn, or wish, without specifying that this is just a *metaphorical* way of speaking; to forget that machines *represent* some mental functions at the level of technology without performing them; and to forget that these deputies of ours act only on command, whether immediate or long-run, is to confuse resemblance with identity, the part with the whole, the form with the essence, thus incurring in magical thinking. Those who write of the *machina ratiocinatrix* may astound the layman, *épater le bourgeois,* or delight the dilettante; but by so doing they hardly deserve to be regarded as the upholders of a tradition of scientific earnestness.

Modern artifacts are marvels of ingenuity, but they are not human and they behave not as humans: if they did, we should not *use* them; artifacts are peculiar physical systems organised by technology to serve man. Is this not enough? Why should we wrap good technology with bad philosophy? And why should the merits of computer designers be attributed to the machines? Why should men "imagine, not only the forms of the gods, but their ways of life to be like their own"? Are there not enough idols without that?