

will be assigned to some independent contractor since the magnitude of the job is about 40 man-months of keypunching, an unfeasible undertaking for the ABF. The results of these various studies will be published as we go along, whenever we have something of significance to report.



REVIEWS

COMPUTERS AND COMMON SENSE

by Mortimer Taube

Columbia University Press (1961)

Nuel D. Belnap, Jr.

Department of Philosophy

Yale University

Can computers be designed and programmed to perform tasks involving “quasi-human intelligence, knowledge, and knowledgeability?” (R1) Taube thinks not. Indeed, in the book under review he labels the pursuit of such a goal a “scientific aberration,” abusing its supporters as like unto the frauds or fools of astrology. (R2)

Taube’s book is on the whole pretty silly; not only does its crimson tone make it a difficult book from which to learn, but it is full of bad arguments purporting to prove that this or that aspect of the program in artificial intelligence is impossible of fulfillment. For antidotes the reader is directed to reviews in *Science* (March 2, 1962) and *Behavioral Sciences* (April, 1962), and to a carefully thought out monograph by Paul Armer, “Attitudes toward intelligent machines,” published by the RAND Corporation.

In this review I should like to concentrate on four connected but distinguishable questions concerning the possibility of developing computers with ability to perform such tasks as mechanical translation of natural languages, playing chess or checkers, simulating the functions of the brain in higher-order problem solving, conducting fully automatic warfare, proving substantial theorems in mathematical logic, writing music or poetry, and the like.

(1) Is it *logically* possible that a machine can be made to do these marvelous things?

(2) Granting that it may be *logically* possible, is it *really* possible; *i.e.*, is there any positive, concrete evidence supporting the claim that machines can be made to do these things?

(3) Granting that it may even be really possible (in the prescribed sense), is it *economically possible*—can we afford it?

will be assigned to some independent contractor since the magnitude of the job is about 40 man-months of keypunching, an unfeasible undertaking for the ABF. The results of these various studies will be published as we go along, whenever we have something of significance to report.



REVIEWS

COMPUTERS AND COMMON SENSE

by Mortimer Taube

Columbia University Press (1961)

Nuel D. Belnap, Jr.

Department of Philosophy

Yale University

Can computers be designed and programmed to perform tasks involving “quasi-human intelligence, knowledge, and knowledgeability?” (R1) Taube thinks not. Indeed, in the book under review he labels the pursuit of such a goal a “scientific aberration,” abusing its supporters as like unto the frauds or fools of astrology. (R2)

Taube’s book is on the whole pretty silly; not only does its crimson tone make it a difficult book from which to learn, but it is full of bad arguments purporting to prove that this or that aspect of the program in artificial intelligence is impossible of fulfillment. For antidotes the reader is directed to reviews in *Science* (March 2, 1962) and *Behavioral Sciences* (April, 1962), and to a carefully thought out monograph by Paul Armer, “Attitudes toward intelligent machines,” published by the RAND Corporation.

In this review I should like to concentrate on four connected but distinguishable questions concerning the possibility of developing computers with ability to perform such tasks as mechanical translation of natural languages, playing chess or checkers, simulating the functions of the brain in higher-order problem solving, conducting fully automatic warfare, proving substantial theorems in mathematical logic, writing music or poetry, and the like.

(1) Is it *logically* possible that a machine can be made to do these marvelous things?

(2) Granting that it may be *logically* possible, is it *really* possible; *i.e.*, is there any positive, concrete evidence supporting the claim that machines can be made to do these things?

(3) Granting that it may even be really possible (in the prescribed sense), is it *economically possible*—can we afford it?

(4) Granting that it may be economically possible, to what extent is this a good way to spend our admittedly limited time and money?

My plan is to take up each of these questions in turn.

(1) Can computers think? If this is meant as a question of *logical* possibility (the "logical 'can'"), the answer will depend—as does any answer to a question of logical possibility—on definitions. If men are (by definition) very complex computers or if computers are (by definition) low grade intellectual organisms, then of course computers can think. Even if men and machines are to be conceived as lying on an n-dimensional continuum as to intelligence (as Armer suggests), the questions as to logical possibility must be given an affirmative answer. But if men have (by definition) essential features which machines fail (by definition) to have—*e.g.*, consciousness or intuition—then obviously the answer must be that machines cannot think.

There is unfortunately some controversy about these definitional matters, though the whole thing smacks of the (*passé*) mechanist-vitalist controversy; Taube is a (*passé*) vitalist. He argues, for example, that machines cannot simulate brain-function, because (by definition) the function of the brain is "to secure successful behavior of the organism in relation to its environment" (p. 71); and machines cannot learn because learning "is the change from conscious to unconscious (habitual) activity to attain a desired goal." (p. 47) Taube pairs off *his* arguments from definition with those of the (*passé*) mechanists: of course machines can do what the brain does, since "nature has given us a working model" [in man] (R3); and of course machines can learn, since chess and checker playing machines improve their game with time, and learning is "improved performance based on consequences of past performance." (p. 47)

This has the ring of word-play. And so it is; but it is not for that reason inconsequential. It is important to realize, as I think the disputants do not, that the argument here is neither "merely" verbal, nor is it an argument about empirical fact; it is rather a dispute about what Kant would have called "regulative principles," principles which are not *results* of inquiry, but which rather *guide* inquiry: regulative principles suggest lines of attack, provide a conceptual frame for asking questions, and even furnish definitions. A man's answer to (1) thus depends on which of the following questions he would answer in the affirmative: (a) Are there "massive analogies" between man and computer? (R4) (b) Or is there a "radical discontinuity" between the two? (c) Or is there, as Armer suggests, a "continuum" stretching between them? To decide between these is to decide between regulative principles. The decision will *show* itself in the kind of experiments a man will think worth doing, in the kind of questions he will ask, and as we here illustrate, in his choice of definitions: if a man subscribes to "massive analogies" as a regulative principle, then he will define "learning" in such a way as to guarantee that it can be truthfully applied to both men and machines; if he subscribes to the principle of "radical discontinuity" he will define "learning" in such a way that a machine cannot meaningfully

(much less truly) be said to learn; and if he chooses the middle of the road "continuum" principle, then he will define "learning" in such a way that its applicability will be a matter of degree. But—and this is the point I should like to make against all disputants—the choice of definitions cannot rationally be used in support of one's choice of regulative principles.

How, then, can one adjudicate between rival regulative principles? It is in this domain, I submit, that the pragmatic theory of truth comes into its own: a regulative principle is "true" if it gives rise to fruitful lines of inquiry, to interesting, answerable questions, and to theoretically fecund schemes of definition.

Taube does in fact argue here and there that the assumption that men are machines (or that computers are organisms) has not led to fruitful and interesting results; that it has given us programs, policies, and prophecies, but has not increased our understanding of either men or machines. More often, however, he is involved either in showing that arguments for the empirical truth of "the myth of thinking machines" (his subtitle) are circular or fallacious (they are—see especially p. 69 and the long note on pp. 132-133), or in advancing equally circular or fallacious arguments for the empirical truth of his own thesis of "radical discontinuity," which we could with propriety call the "myth of consciousness." (R5) His chief weakness is his total failure to support the principle of "radical discontinuity" as a *regulative* principle. Nowhere, except perhaps in his muddy chapter on meaning, does he indicate how acceptance of his principle would lead to interesting questions or viable lines of research. He wants us to reject the programs growing out of the "massive analogies" or "continuum" principle, but gives us not the smallest hint of a program which could take their place.

(2) In order even to *ask* the second question, "Is there any positive evidence giving us reason to suppose that computers can be built and programmed to translate languages, play good chess, prove substantial mathematical theorems, etc.?" we must adopt the "massive-analogy" or the "continuum" framework, so that it is at least logically possible that a machine could do these things. But—and here Taube argues cogently—it is important to realize that merely allowing the logical possibility does not answer the question as to the *real* possibility, unless one adopts the ludicrous (and common) course of arguing from the premiss of mere logical possibility, or from the premiss that there is no evidence against artificial intelligence, to the conclusion that there is some positive evidence for it. Having usefully disposed of this *argumentum ad ignorantiam*, Taube goes on to detail his claim that in fact there is no evidence supporting the familiar prophecy (Simon and Newell) that "within ten years computers will discover important mathematical theorems, write worthwhile music, take over the major part of the field of psychology, and dethrone the current world chess champion" (p. 128), or that mechanical translation is just around the corner.

Taube's arguments for the real impossibility of mechanical trans-

lation (taking that as a test case) are uniformly silly. Some of them tacitly seem to suppose that what is wanted is a perfect mechanical translator (p. 35)—as if any human were a perfect translator! Some rely on misapplications of Gödel's incompleteness theorem, partly arising from a confusion of syntax and semantics (*e.g.*, p. 34, p. 57): Gödel showed only that we cannot formalize the notion of mathematical *truth*, not at all that we cannot formalize mathematical grammaticalness (which indeed he formalizes). Some of Taube's arguments rely on the claim that there is present in humans a faculty of "intuition" or "consciousness" denied to machines (pp. 34-35, 93, 96)—this is but a pseudo-empirical restatement of his regulative principle of "radical discontinuity." One of his arguments has as a premiss the assertion that for formal translatability between two languages A and B, there must be "at some level a one-to-one relation of synonymy between elements of A and elements of B." (p. 33) This last is demonstrably false, inasmuch as there are many examples of effectively inter-translatable formal languages, where in neither direction is the translation-function one-to-one. And as a last example, he in one place argues from the ambiguous philosophical premiss that the minimum carrier of meaning is the sentence, to the remarkable conclusion that the dictionary for any adequate translation program must contain every English sentence as a separate entry (p. 29; see also p. 36).

These are all general arguments, different in kind from arguments based on individual computer accomplishments such as the chess-playing machine or the theorem-prover. The obstacle to clear thinking about evidence of this latter sort is that it is difficult to say just what ought to count as "evidence" for the possibility of artificial intelligence on the part of computers, when on the one hand exaggerated claims by certain excitable members of the computing industry get compounded with wildly anthropomorphic descriptions of actual or proposed computer set-ups (*R6*); and when on the other hand anything that any computer has actually done gets derogated by others as "purely mechanical" and just not good enough to point toward the possibility of artificial intelligence. Some clear-headed analytical work needs to be done on this question of evidence; however, even if we dismiss as irrelevant all computer accomplishments to date, there remains one argument in favor of the real possibility of artificial intelligence: whenever in the past we have set a lot of bright people to work trying to solve a difficult problem, we have generally found a solution—and this fact is a positive reason in support of the view that we can do so again.

(3) This brings us to the third question, that concerning economic possibility: can we afford to turn loose on this problem sufficiently many bright people for a sufficient length of time? As far as I know, *nobody* has made any adequately grounded estimates of the cost of developing (say) an economically useful mechanical translator capable (say) of reading from a wide range of printed texts. Nor does anyone know (as far as I am aware) how much it would cost to build a world-champion chess