Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Book Reviews

MANAGEMENT SCIENCE
Vol. 11, No. 1, September, 1964
Printed in U.S.A.

BOOK REVIEWS

FEIGENBAUM, EDWARD AND FELDMAN, JULIAN (eds.), *Computers and thought.* New York: McGraw-Hill, 1963, $7.95, 535 pp.

This work is remarkable in that although it is composed of contributions by twenty-eight authors, it reads not like a typical multiple author compendium but like a real book on a well defined subject. This bespeaks for the emergence of Artificial Intelligence as a mature field with well established directions, methods, theoretical concepts, and, above all, clearly delineated proximal and distal goals.

Another sign of rapid maturation of the field, barely in its second or third decade (depending on whether one reckons from the first systematic mechanization of logical processes or from the first electronic computers) is the natural connection that has been established between the emerging philosophical and technological problems.

The science of artificial intelligence certainly does not "solve" the philosophical question to which it has given rise, namely "Do machines think?" But then no philosophical problem has ever been solved except in the estimation of those who cling to dogmas. Scientists, not philosophers, "solve" problems, and they are not obligated to solve problems as they are posed by the philosophers. However, neither is it proper to dismiss the problems posed by philosophers, something all too easy to do on the grounds that the questions are so badly put that their meaning disintegrates under the slightest pressure of analysis. The relation between philosophy and science ought to be that between a challenge and the acceptance of a challenge. This relation has apparently been achieved with respect to the philosophical questions instigated by computer technology, such as "What is intelligence?" and even "What is thought?"

A discussion of these matters opens the volume under review, namely a paper by Turing, one of the founders of the new science. Turing's argument, defending the affirmative answer to the question whether machines can think, proceeds by a successive refutation of the usual objections against ascribing the thinking function to the machine. As is usual in discussions of this sort, the importance of the arguments is not in whether the question has been convincingly answered one way or another but rather in the issues raised in the process of analysis.

Those who would deny "thinking" to machines play a part in the science of artificial intelligence analogous to that of the vitalists in biology. Their role is to say "Yes, but . . . ." When organic matter was shown to be synthesizable, the vitalists said "Yes, but living beings are autonomous sources of energy." When the law of conservation of energy was shown to apply also to living beings, the vitalists said "Yes, but living beings are guided by teleological, not mechanistic laws." When cyberneticians showed that the distinction between teleological and mechanistic laws was an artificial one, the vitalists retreated to theological positions. The importance of this development was not in who won the argument but in the concomitant deepening analysis of the life process.

A very similar retreat is being forced on the "negativists", those who would at all costs deny that "machines think." Again the outcome is not a victory for one side or the other but the concomitant excursion into the problem areas related to the nature of the thought process. For with each "victory," the nature of the inadequacy of mechanized thought became progressively clearer and with this realization comes the recognition of the important problem areas to be still explored in the preparation of a new "victory." In this way, the rarely achieved union of philosophy and technology, of introspection and experiment has been continued.

The first "thinking machines" performed only arithmetical operations. These machines were oversold to the public by means of sensational metaphors and science fiction type extrapolations. In the reaction that followed, the chief emphasis was put on the mechanical, blindly obedient, "brilliantly idiotic" aspects of computer operations. Although before the advent of the electronic computers, routine calculating feats might have been considered as prima facie evidence of "thought," they ceased to be so considered after they were actually performed by devices to which it was important for various reasons to deny the faculty of intelligence. Accordingly, a retreat was made in the definition of thought. "Intelligent choice" was now seen to be the crucial ingredient of genuine thought. But it soon became apparent that whenever "intelligent choice" could be formalized, it could also be mechanized, at least in principle, by providing a way to compare all possible outcomes of alternative decisions with a pre-assigned hierarchy of values. Indeed, the fundamental results in the theory of games made it clear that *in principle* a rational play of a (zero-sum) game is a matter of choosing from a finite set of strategies a single one or a frequency distribution over the strategies, characterized by a specifiable mathematical property. A machine which played a perfect game of tic-tac-toe was exhibited, if I recall correctly, in the early 1950's.

However, the phrase "in principle" became the keystone of the next and the currently crucial argument of the negativists, for the method of exhaustive search, employed by the tic-tac-toe machine, becomes useless in the game of checkers, let alone chess. To apply the method in chess, some 10^{120} continuations must be examined. Since less than 10^{16} microseconds are available in a century, to "solve" chess by this method would take more centuries than there are atoms in the universe, assuming a fantastic speed of examining at the rate of a million continuations per second.

Obviously a mechanization of chess strategy by the method of exhaustive search is beyond physical possibility. How then do thinking beings play chess? At present this question can be answered only by introspection. Chess players agree on certain by no means rigorously formalizable, yet intuitively understandable "holistic" principles of chess strategy. These are coupled with *partially* exhaustive search procedures along *selected* continuations. It seems, then, that the recognition and evaluation of certain *patterns* and the selection of certain "promising" lines of search constitute intelligent strategic thought as distinguished from mechanized procedures. To the extent that the challenging objections of the negativists have intellectual content (as opposed to emotional revulsion, of which I shall speak below) this is the very kernel of the issue presently alive. Thought is felt to be coupled with the perception of wholes, with hunches and insights, with the ability to conduct searches not mechanically (step by step) nor only with regard to ultimate rigidly fixed goals but also with regard to sub-goals, often invented in the process of the search itself. These will be recognized as the components not only of strategic thought but of all creative intellectual endeavor.

Computers and Thought is devoted to precisely these problems. The fundamental axiom of cybernetics[1], "What can be formalized can be programmed" is obviously true in the converse, "What is to be programmed must be formalized." The direct problem, that of programming what has already been formalized, is a purely technical problem. It is the inverse problem, namely that of formalizing what is perceived and learned in the process of intuitive understanding and in genuine thought which constitutes the focal point of creative effort in cybernetics today. It is this unity of theme which welds the volume compiled by Feigenbaum and Feldman into a genuine book.

The opening essay (Turing, I, 1)[2], entitled "Can machines think?" poses the philosophical problem already mentioned and, in answering the standard arguments of the negativists, lays bare the problems with which modern creative cybernetics is concerned.

The next three sections deal with problem-solving machines, which operate on the heuristic principle. This is the principle of artificial intelligence which is replacing the method of exhaustive search. In chess, for example, a "heuristic" involves a formalization of some general principles recognized as important by chess players. Some of these are easy to formalize. For example, the relative values of the pieces, (Q = 9, R = 5, etc.) have been long established. Accordingly, forthcoming exchanges can be evaluated in these terms. Other principles, e.g., "mobility" are more difficult to formalize. Yet this too can be done, for example, in terms of the number of open alternatives associated with a position, etc. A heuristic assigns relative values to the various components (King's safety, material, mobility, etc.) Once evaluations have been assigned, the minimax principle (which assumes that the opponent will choose replies best for *him*) can be applied to the choice of the next move. (Newell, Shaw, and Simon, I, 2).

The same sort of principles apply in proving theorems. The most formalized

[1] In the U.S.S.R. "cybernetics" is used to refer not only to the science of regulation and control, as with us, but also to the entire field comprising artificial intelligence, computer technology, and its applications. I believe this extension of usage is justified and ought to be generally adopted in view of the powerful abstract and general ideas (e.g., information theory) which unify those various fields. It seems that the scope of cybernetics is as well defined as that of physics. If matter and energy comprise the fundamental subject matter of the latter, organization and information comprise the subject matter of the former. The notion of entropy is the link between the two, as has been shown by Szilard, Brillouin, and others.

[2] I shall refer to the authors of the articles by names in parentheses followed by the part and section where they appear.

mathematical system is that of mathematical logic embodied in the *Principia* of Whitehead and Russell. It is possible to derive any proposition of that system from the basic axioms by applying very few transformation rules, by which strings of symbols are replaced by other strings of symbols. This is mathematics in its "abstractest", and, according to one philosophy of mathematics, all mathematics can be made to fit into this scheme. In principle, therefore, it ought to be possible to prove all mathematical propositions or, at least all the propositions in the *Principia* by the method of exhaustive search. But as in chess, the method is out of the question practically speaking. Accordingly, heuristics have to be designed. These involve sub-goals. A man faced with the problem of proving a mathematical theorem might think so; If I can't prove this proposition directly, maybe I can prove another one, from which the assigned proposition is clearly a consequence." In proving theorems of plane geometry (Gelernter, I, 3; Gelernter, Hansen and Loveland, I, 3) heuristics involve "drawing auxiliary features in the diagrams." (Actually already the use of a diagram involves a heursitic.) A man draws such additional features (e.g., connects points by line segments) in order to help his logical intuition by visual perception. The analogue in the theorem-proving machine is the creation of additional features to refer to in the proof. For example propositions about parallelograms can be easily proved by reference to triangles into which the parallelogram can be divided by diagonals.

At this point the negativist can still hold on to his argument. The "heuristics," he can point out, have been programmed into the computer and so are, in principle, no different from detailed instructions mechanically followed. The rejoinder to this may well be that heuristics are also "programmed" into intelligent beings. People learn how to play chess and how to prove theorems from other people who tell them what to look for and what strategems to employ. To this the negativist has two additional rejoinders. One is that human beings are taught not by having their neuronal connections specifically connected but by the use of language. There is still no machine which can "understand" human language, i.e., understand the *intent* of the speaker in spite of the vagueness, the ellipsis, and the metaphorical circumlocutions inherent in human communication. The other rejoinder is that in any line of endeavor human beings can improve their performance and that this improvement is by no means limited to learning perfectly what one is taught but can carry the student beyond the competence of the teacher.

It turns out that the second objection is easier to meet than the first. Learning behavior is comparatively easy to formalize and hence to simulate in machines. Techniques of this sort are discussed in some of the papers (Uhr and Vossler, I, 6; Feigenbaum, II, 2) and have also been developed in other contexts ("learning" in initially randomly connected simulated neural nets, not discussed in the volume). Moreover, "learning machines" can at times surpass their "teachers" as in the case of the checker playing machine (Samuel, I, 2) which has beaten its creator.

This brings us to the most "creative" efforts in modern cybernetics. One standing problem is that of programming a machine to understand human language. Such understanding, be it noted, is a matter of degree. Even a desk calculator "understands" simple commands, such as "add", "multiply", etc. Elaborate computers "understand" much more elaborate languages (e.g., FORTRAN). Advance in this field is effected by making the languages which can be "understood" by the machine more and more similar to natural languages. This involves the problem of constructing a rigorous grammar of a natural language and connects cybernetics to structural linguistics. The topic is touched upon by Lindsay (I, 5). The problem of constructing a rigorous and complete grammar of a natural language has remained largely unsolved and with it the vastly more difficult problem of constructing a rigorous semantics of a natural language. Here the practitioners have followed research strategies analogous to those developed in automated problem-solving, namely circumscribing a small area and learning the "art" from the emergent problems. A good example is given by the "automatic question answerer on the subject of baseball" (Green, Wolf, C. Chomsky, and Laughery, I, 5).

In this connection, it should be noted that trivial "automatic question answerers" are found in every amusement park, namely, the fortune telling machines which "answer" pre-set questions like "Does she love me?" and "Will I make a lot of money?" These machines have the same relation to the baseball question answerer as the pre-set answer calculating machine has to the computer. A pre-set answer calculating machine would be one in which the answer to every problem is pre-set specifically. For example, punching buttons labeled "1", "3" and "add" trips off the one relay which makes "4" show in the output. Such a machine does not exist; it would be silly to construct one. A calculator able to add sums (pairs only) up to 10,000 would require fifty million connections. Calculators, of course, do not work in this way. Even a desk calculator is programmed with the *rules* of addition, which enable it to sum arbitary sets of numbers to the limit of its capacity. The same is true of the automatic question answerer. It is programmed to transform arbitrary questions (within the context of a certain vocabulary and certain syntactical rules) into information retrieval processes, i.e., calculations performed on the data (baseball statistics in a certain year) to answer the questions within its scope of "understanding".

Already here the grammatical-semantic ambiguity of natural language makes itself felt. Consider, for instance, the question "Did the Red Sox win most of their games in July?" This could mean "Is it true that most of the games which the Red Sox played in July were won by them?" or "Is it true that most of the games won by the Red Sox were games played in July?" or "Is it true that the Red Sox won more games in July than in any other month?" The authors say that the question has five different meanings. I could not think of the other two.

Understanding natural language is imbedded in the more general problem of pattern recognition. This problem is specifically represented by two papers of the volume (Lindsay, I, 5; Feigenbaum, II, 2). Actually it is inherent also in most of the other papers concerned with problem-solving, since the problem emerges as soon as the method of exhaustive search (or specifically programmed access to questions) is abandoned. Herein is also the challenge of simulating truly human "intelligence," for there is no question that this intelligence resides in the symbolic function of language, which "slices up" the universe for us into verbal categories and relegates our thinking to the manipulation of these categories apart from concrete experience. Herein resides the blessing conferred upon the human race and, incidentally, also the curse, for our compulsions, superstitions, and suicidal conflicts owe as much to our virtuosity with symbols as do our science, moral sense, and art.

Computer simulation has at times been grossly misinterpreted in popular estimation as the "automation of decision-making", which conjures up visions of a world ruled by mechanical brains. The content of Part II, devoted to the simulation of cognitive processes ought to disspell much of this misapprehension through its emphasis on the true importance of this method as an adjunct to behavioral theory. A simulation program is nothing but a way of calculating the consequences of a particular mathematical model (i.e., a set of strictly formalized assumptions) of some process. If the model is sufficiently simple, its consequences can be derived by analytical methods. Comparison of the consequences of the model with appropriate observations constitutes a test of the model (i.e., of the assumptions). If the model is so complex that the derivation cannot be accomplished analytically, one can abstract particular numerical results by a simulation program. The possibility of doing so (since computers have become available to behavioral scientists) has stimulated the construction of mathematical models of situations hitherto ignored because of their complexity. In this way, the range of application of mathematics to behavioral theory has been greatly expanded thanks to computers.

A computer-simulated model of elementary social behavior (Gullahorn and Gullahorn, II, 4) is a typical example. The authors take some propositions, advanced by some social psychologists in verbal form about interactions of people in small groups, sharpen the propositions by quantifying the variables, and simulate the process to observe the "results". These can be directly compared with experimental data (also appropriately quantified) or, usually with greater success, can be used to answer the following sort of question: If people (or other units) indeed interacted as the social psychologists (or economists or political scientists) say they do, what would the results be like? No restrictions need be placed on the quality of the interactions. They can be highly "rational" as postulated in game theory, or quite irrational, governed, say, by immediate hedonistic responses to the situation. In short, simulation provides the opportunity of experimenting in the realm of the possible, when experimentation in the realm of the actual is difficult.

Little can be learned by this method about how people actually act. For the most part, the models are highly over-simplified idealizations of social situations (Feldman, II, 3). Moreover, it is possible that the acquisition and dissemination of whatever knowledge is obtained in simulations will serve to change the situations simulated. A case in point is the simulation of the trust investment process (Clarkson, II, 3). If all the investors were guided by the results of the simulation, the whole character of the market might change and hence the investment process might also change. However, the lack of "realism" or practical usefulness of such models matters not at all if the purpose of simulation is well understood—not the prediction of human behavior in specific instances but bringing into focus some characteristic and general features of behavior (in the present case of the cognitive process) is the goal of simulation as an adjunct to behavioral theory. This spirit pervades also the remaining papers in Part II (Newell and Simon, II, 1; Feigenbaum, II, 2; Hunt and Hovland, II, 2).

Part III, Survey of Approaches and Attitudes, contains two papers. That of Armer (on attitudes toward artificial intelligence) overlaps to a considerable extent with Turing's introductory paper, not necessarily a shortcoming, in view of the way it serves in a summing up capacity. That of Minsky is an excellent survey of the state of the art to within a year or two of the date of publication. In this paper many of the points emphasized in this review (and some others) are brought out with force and clarity.

Part IV is a "selected descriptor-indexed bibliography to the literature on artificial intelligence" also by Minsky.

The main virtue of the book namely, its remarkable cohesion, has already been mentioned. Another major virtue is the consistent clarity and informality of style, which makes the book unquestionably accessible to the non-specialist. It is a book devoted to ideas, not to details of engineering or to intricacies of special techniques. Moreover, the ideas are, or ought to be, of deep and immediate concern to all thinking people.

Much of the book (though not its principal theses) will probably become dated in a very few years, but this is something that could not be helped by either the authors or the editors. The only reservation which occurs to me is the failure on the part of both Turing and Armer, who both polemicize against the negativists, to bring out the one valid basis of the negativists' concern. This concern is obscured, not clarified, in the "theological arguments" and in the "emotional outbursts" against anti-utopian visions.

The valid basis of the negativists' attitude is, I believe, a frankly religious one, for it is related to the sin of idolatry. But it need not on that account be relegated to the "theological" level, where an argument which makes sense to a secular mind cannot be made. The sin of idolatry, on the other hand, can be exemplified in a simple, concrete situation. Suppose that computers were physically dangerous, as many engines are, and suppose a man got caught in the gears, so that the only way to save him would be to smash the computer. I do not think that there would be much hesitation about what to do. Nor would any markedly utilitarian considerations enter the decision. A man is (ought to be) "worth" more than the computer even if he is old, feeble, and demented, and even if the computer is worth several million dollars and is the proudest creation of human genius. The decision to save the man at the cost of the computer can only stem from our *willingness* to assume a "consciousness" in the man and our *unwillingness* to assume the same in the computer, no matter what "experimental" evidence there may be to the contrary. To do the opposite would be to commit the sin of idolatry. The extension of consciousness, thought, feelings, etc. to other human beings is not a rational act. It cannot be defended by either induction or deduction any more than the refusal to extend these faculties to organized hardware can be defended by rational argument, as both Turing and Armer have conclusively shown and as this reviewer also believes.

However, life would be impossible (or at least unbearable) if we did not make *some* commitments on other than rational grounds. Most people would say that it is not permissible to kill a baby in order to extract a substance with which to cure one thousand cancer patients. This conviction cannot be defended by any utilitarian calculus except in a circular manner. An argument is sometimes offered to the effect that a society which allows human sacrifice is not likely to be viable, but such an argument is not conclusive, and even if it were, the low viability of such a society might well stem from its failure to command the loyalty of its members and so again reduces to a nonrational commitment. The only sound argument against human sacrifice is an emotional argument; human sacrifice is wrong. Although it behooves a rational man to avoid such arguments as long as possible, i.e., to reduce moral issues to rational ones to the fullest *possible* extent, I think that ultimately some non-rational commitments are unavoidable, and moreover I do not think that these commitments can all be reduced to a single, simple one, such as "The *summum bonum* is to maximize the probability of survival" (as some neo-utilitarians would have it). No one has ever calculated such probabilities nor has stated the assumptions on which such calculations could be based.

It appears, then, that the only possible defensible ground on which the negativists could make a stand is the admittedly non-rational (i.e., religious) commitment against idolatry. If one fears that "thinking" gives the computer a claim on human empathy and if one fears that the extension of empathy to computers may jeopardize the extension of empathy to men (or to living beings, or to God, as the believers would have it), then attributing thought to computers can indeed be viewed as idolatry.

In our day idolatry is very real and, in the opinion of serious thinkers (for example Norbert Wiener, with whose name cybernetics will be forever associated), very evil. We worship all kinds of things to the neglect of man. To all too many, the target-seeking missile is a thing of beauty. I am not sure that such matters should necessarily have been discussed in *Computers and Thought*. However, since the "theological issue" *has* been raised by Turing (I do not know how seriously), perhaps the genuinely religious or the humanistic issue should have been raised also, at least in passing.

ANATOL RAPOPORT
University of Michigan

EWING, DAVID W. (ed.). *Long-range planning for management.* New York City: Harper & Row, Publishers, 1964, $9.75, 566 pp.

At first glance this book may seem to have the characteristics of an "Anthology", not in the original Greek meaning of flower gathering, referring to a collection of poems, but in its modern meaning of a collection of writings of miscellaneous authorship to present an idea.

David Ewing, who is associate editor of the Harvard Business Review, has done more than select the most appropriate papers from the available literature on long-range planning; through his comments and conclusions in the introductory chapter to each major section, and particularly through the arrangement of the subject, he presents a valuable and most readable book for the executive as well as the student.

He has arranged the subject of long-range planning into the following major sections: Nature and Principle, Strategy, Organization, Steps in Making a Plan, Special Functions and Areas, Conclusion. Each Section consists of five to fifteen individual chapters by different authors.

The new edition has doubled in size from the 1958 edition with more emphasis on strategy and actual organizational experiences. Among the contributors are: Mark Cresap, Jr., "Some Guides and Standards"; Peter Drucker, "Long-Range Planning Means Risk-Taking" and "The Divorce of Planning from Doing"; Harry L. Hansen, "The Timing of New Product Introduction"; Bruce Payne, "Steps in Long Range Planning"; James Quinn, "Top-Management Guides for Research Planning"; L. Eugene Root and George Steiner, "The Lockheed Aircraft Corporation Master Plan"; George Steiner, "Stages in Development"; Melvin Salveson, "The Strategy of Innovation in Technical Industries"; Stewart Thompson, "Organization in a Steel Industry"; and forty other authors.

David Ewing's introduction to each major section is not only a contribution to the subject of long-range planning by a critical discussion of the specific aspects but also a guide for the reader toward the chapters of particular interest. Brief biographical notes on the authors are included with the introduction for each section.

While in 1958, at the time of the first edition, less than two hundred corporations had formal long-range planning departments, over seven hundred have such departments today and many smaller companies are involved in long-range planning through part-time assignments to executives.

The decline and demise of many a company can be traced, at least partially, to the neglect of true corporate or long range planning. Executives in top management who are too busy to scan for papers on long-range planning scattered throughout various publications will find the format of a book more appealing for reference than reprints or clippings from journals.

The newcomer to the task of long-range planning as well as the seasoned executive should find many valuable thoughts in this book that will aid them in thinking through long-range planning problems; they will also be interested in some of the actual experiences of organizations reported in this book, although they obviously cannot be considered ready models to be copied.

David Ewing's closing chapter discusses "Knowledge of Planning Methods," "Knowledge of Conditions and Trends Affecting Planning," "Knowledge of Desirable Goals and Objectives." In his "Needs For The Future" he makes an appeal for release of more concrete cases and experiences, as well as more penetrating research.

Leopold R. Michel
Marblehead, Massachusetts