Reprinted from OPERATIONS RESEARCH
Vol. 6, No. 1, Jan.-Feb., 1958
Printed in U.S.A.

HEURISTIC PROBLEM SOLVING: THE NEXT ADVANCE IN OPERATIONS RESEARCH*

Herbert A. Simon and Allen Newell

Carnegie Institute of Technology, Pittsburgh, Pennsylvania, and
The Rand Corporation, Santa Monica, California

THE IDEA THAT the development of science and its application to human affairs often requires the cooperation of many disciplines and professions will not surprise the members of this audience. Operations research and management science are young professions that are only now beginning to develop their own programs of training; and they have meanwhile drawn their practitioners from the whole spectrum of intellectual disciplines. We are mathematicians, physical scientists, biologists, statisticians, economists, and political scientists.

In some ways it is a very new idea to draw upon the techniques and fundamental knowledge of these fields in order to improve the everyday operation of administrative organizations. The terms 'operations research' and 'management science' have evolved in the past fifteen years, as have the organized activities associated with them. But of course, our professional activity, the application of intelligence in a systematic way to administration, has a history that extends much farther into the past. One of its obvious antecedents is the scientific management movement fathered by FREDERICK W. TAYLOR.

But for an appropriate patron saint for our profession, we can most appropriately look back a full half century before Taylor to the remarkable figure of CHARLES BABBAGE. Perhaps more than any man since Leonardo da Vinci he exemplified in his life and work the powerful ways in which

* Address at the banquet of the Twelfth National Meeting of the OPERATIONS RESEARCH SOCIETY OF AMERICA, Pittsburgh, Pennsylvania, November 14, 1957. Mr. Simon presented the paper; its content is a joint product of the authors. In this, they rely on the precedent of Genesis 27:22, "The voice is Jacob's voice, but the hands are the hands of Esau."

fundamental science could contribute to practical affairs, and practical affairs to science. He was one of the strongest mathematicians of his generation, but he devoted his career to the improvement of manufacturing arts, and—most remarkable of all—to the invention of the digital computer in something very close to its modern form.

The spirit of the operations researcher, his curiosity, his impatience with inefficiency in any aspect of human affairs, shows forth from every page of Babbage's writing. I give you just one example:

> Clocks occupy a very high place amongst instruments by means of which human time is economized: and their multiplication in conspicuous places in large towns is attended with many advantages. Their position, nevertheless, in London, is often very ill chosen; and the usual place, half-way up on a high steeple, in the midst of narrow streets, in a crowded city, is very unfavourable, unless the church happen to stand out from the houses which form the street. The most eligible situation for a clock is, that it should project considerably into the street at some elevation, with a dial-plate on each side, like that which belonged to the old church of St. Dunstan, in Fleet-street, so that passengers in both directions would have their attention directed to the hour.[1]

I have mentioned Babbage as the inventor of the computer. Since Babbage and the computer are going to be the heroes of my talk tonight, I should like to tell you a true story, culled from Babbage's writings, about the history of the computer. I like this story because it illustrates not only my earlier point about the many mutual relations of the professions in our field, but also because it gives the underdogs like myself—trained in 'soft' fields like economics and political science—something we can point to when the superior accomplishments of the natural sciences become too embarrassing for us. As you will see, this story shows that physicists and electrical engineers had little to do with the invention of the digital computer—that the real inventor was the economist Adam Smith, whose idea was translated into hardware through successive stages of development by two mathematicians, Prony and Babbage. (I should perhaps mention that the developers owed a debt also to the French weavers and mechanics responsible for the Jacquard loom, and consequently for the punched card.)

The story comes from a French document, which Babbage reproduces in the original language. I give it here in translation:

> Here is the anecdote: M. de Prony was employed by a government committee to construct, for the decimal graduation of the circle, logarithmic and trigonometric tables which would not only leave nothing to be desired from the standpoint of accuracy, but which would constitute the most vast and imposing monument of calculation that had ever been executed or even conceived. The logarithms from 1 to 200,000 are a necessary and essential supplement to this work. It was easy for M. Prony to convince himself that even if

> he associated with himself three or four experienced collaborators the longest reasonable expectation of the duration of his life would not suffice to complete the undertaking. He was preoccupied with this unhappy thought when, finding himself before a bookstore, he saw the beautiful edition of Adam Smith published in London in 1776. He opened the book at random and chanced upon the first chapter, which treats of the division of labor and where the manufacture of pins is cited as example.
>
> Hardly had he perused the first pages when, by a stroke of inspiration he conceived the expedient of putting his logarithms into production like pins. He was giving, at this time, at the Ecole Polytechniques, some lectures on a topic in analysis related to this kind of work—the method of differences and its applications to interpolation. He went to spend some time in the country and returned to Paris with the plan of manufacture that has been followed in the execution. He organized two workshops which performed the same calculations separately, and served as reciprocal checks.[2]

It was Prony's mass production of the mathematical tables, in turn, that suggested to Babbage that machinery could replace human labor in the clerical phases of the task, and that started him on the undertaking of designing and constructing an automatic calculating engine. Although the complete absence of electrical and electronic components, and his consequent dependence on mechanical devices, robbed him of full success in the undertaking, there is no doubt that he understood and invented the digital computer—including the critically important idea of a conditional transfer operation.

It would be hard to imagine a more appropriate illustration of the unexpected ways in which human knowledge develops, and of the contribution of all the sciences and arts to this development that is so characteristic of operations research and management science.

As WE TURN our gaze now from past to future, I should like to outline my main thesis quite bluntly. Operations research has made large contributions to those management decisions that can be reduced to systematic computational routines. To date, comparable progress has not been made in applying scientific techniques to the judgmental decisions that cannot be so reduced. Research of the past three years into the nature of complex information processes in general, and human judgmental or heuristic thinking processes in particular, is about to change this state of affairs radically. We are now poised for a great advance that will bring the digital computer and the tools of mathematics and the behavioral sciences to bear on the very core of managerial activity—on the exercise of judgment and intuition; on the processes of making complex decisions.

Let me spell out this thesis, first describing the present situation in operations research as I see it, then indicating why I think this situation is going to change drastically.

THE RAPID GROWTH of operations research over the past two decades has brought to industry and government an important kit of tools for grappling with the complexities of managing large organizations. These tools have been collected from the far corners of the intellectual world—from mathematics, from statistics and probability theory, from econometrics, from electrical engineering, and even from biology. Such exotic techniques as linear programming, queuing theory, servomechanism theory, game theory, dynamic programming, marginal analysis, the calculus of variations, and information theory are now at work helping to solve practical problems of business operation.

Skeptical—and sensibly skeptical—managements have come to see that, even if not all the blue-sky claims for the new approaches have been backed by solid fact, there is a large core of valid technique and application. The tools have produced tangible results in a substantial number of demonstration installations, and the question is less and less 'Are they here to stay?' and more and more 'How and where can we use them effectively?' The traditional areas of production and inventory control, of scheduling, and of marketing research are undergoing a substantial and rapid evolution.

Having observed this important change, we can note with equal accuracy that large areas of managerial activity—it would be correct to say most areas—have hardly been touched by operations research or the advances in management science. Operations research has demonstrated its effectiveness in dealing with the kinds of management problems that we might call 'well structured,' but it has left pretty much untouched the remaining, 'ill structured,' problems.

The trouble, as executives are fond of pointing out to operations researchers, is that there are no known formal techniques for finding answers to most of the important top-level management problems. Nor do these problems seem to be of the same kind as the more tangible middle-management situations in which existing operations research techniques have been most effective. Unarmed with formal techniques, operations researchers have to resort to the same common sense and human cleverness that has served managements these many years. Executives still find a vast sphere of activity in which they are secure from the depredations of mathematicians and computers.

Let me try to make a little more precise this distinction between well-structured and ill-structured problems that today establishes the jurisdictional boundary beyond which formal tools do not reach.

A problem is well structured to the extent that it satisfies the following criteria:

1. It can be described in terms of numerical variables, scalar and vector quantities.

2. The goals to be attained can be specified in terms of a well-defined objective function—for example, the maximization of profit or the minimization of cost.

3. There exist computational routines (*algorithms*) that permit the solution to be found and stated in actual numerical terms. Common examples of such algorithms, which have played an important role in operations research, are maximization procedures in the calculus and calculus of variations, linear-programming algorithms like the stepping-stone and simplex methods, Monte Carlo techniques, and so on.

In short, well-structured problems are those that can be formulated explicitly and quantitatively, and that can then be solved by known and feasible computational techniques.

What, then, are ill-structured problems? Problems are ill-structured when they are not well-structured. In some cases, for example, the essential variables are not numerical at all, but symbolic or verbal. An executive who is drafting a sick-leave policy is searching for words, not numbers. Second, there are many important situations in everyday life where the objective function, the goal, is vague and nonquantitative. How, for example, do we evaluate the quality of an educational system or the effectiveness of a public relations department? Third, there are many practical problems—it would be accurate to say 'most practical problems'—for which computational algorithms simply are not available.

If we face the facts of organizational life, we are forced to admit that the majority of decisions that executives face every day—and certainly a majority of the very most important decisions—lie much closer to the ill-structured than to the well-structured end of the spectrum. And yet, operations research and management science, for all their solid contributions to management, have not yet made much headway in the area of ill-structured problems. These are still almost exclusively the province of the experienced manager with his 'judgment and intuition.' The basic decisions about the design of organization structures are still made by judgment rather than science; business policy at top-management levels is still more often a matter of hunch than of calculation. Operations research has had more to do with the factory manager and the production-scheduling clerk than it has with the vice-president and the Board of Directors.

I am not unaware that operations researchers are often called in to advise management at top levels and regarding problems of the kinds I have called ill-structured. But I think we all recognize that when we are asked by management to advise on such decisions, we are asked because we are thought to possess a certain amount of experience and common sense, and not because of any belief that our specialized tools, mathematical or otherwise, have much to do with the task at hand. I think most of us can distinguish pretty clearly between the cases in which we are working as operations researchers, and those in which we are performing as general management consultants. And I am sure that most of us look forward to the day when our science will enable us to handle with appropriate analytic tools those problems that we now tackle with judgment and guess.

The basic fact we have to recognize is that no matter how strongly we wish to treat problems with the tools our science provides us, we can only do so when the situations that confront us lie in the area to which the tools apply. Techniques are the arms and hands of science, and the reach of a science is measured by their range. The telescope made sunspots and Jupiter's moons a part of Galileo's science, just as particle accelerators and the mathematical machinery of quantum mechanics bring the interior of the atom within the reach of the nuclear physicist.

In dealing with the ill-structured problems of management we have not had the mathematical tools we have needed—we have not had 'judgment mechanics' to match quantum mechanics. We have not had the engines—no executive centrifuges. We have had only the rudiments of experimental techniques for observing organizational behavior in the laboratory, although we have made great strides in the last decade in developing these.

IF OUR SCIENCE, then, is to be coextensive with the field of management, we must have the tools and techniques that will extend its range to that whole field. I think there is good reason to believe that we are acquiring these tools and techniques at this very point in history.

Even while operations research is solving well-structured problems, fundamental research is dissolving the mystery of how humans solve ill-structured problems. Moreover, we have begun to learn how to use computers to solve these problems, where we do not have systematic and efficient computational algorithms. And we now know, at least in a limited area, not only how to program computers to perform such problem-solving activities successfully; we know also how to program computers to *learn* to do these things.

In short, we now have the elements of a theory of heuristic (as contrasted with algorithmic) problem solving; and we can use this theory both to understand human heuristic processes and to simulate such processes with digital computers. Intuition, insight, and learning are no longer exclusive possessions of humans: any large high-speed computer can be programmed to exhibit them also.

I cannot give here the detailed evidence on which these assertions—and very strong assertions they are—are based. I must warn you that examples of successful computer programs for heuristic problem solving are still very few. One pioneering effort was a program written by O. G.

SELFRIDGE and G. P. DINNEEN that permitted a computer to learn to distinguish between figures representing the letter *O* and figures representing *A* presented to it 'visually.'[3] The program that has been described most completely in the literature gives a computer the ability to discover proofs for mathematical theorems—not to verify proofs, it should be noted, for a simple algorithm could be devised for that, but to perform the 'creative' and 'intuitive' activities of a scientist seeking the proof of a theorem. The program is also being used to predict the behavior of humans when solving such problems. This program is the product of work carried on jointly at the Carnegie Institute of Technology and the Rand Corporation, by Allen Newell, J. C. Shaw, and myself.[4]

A number of investigations in the same general direction—involving such human activities as language translation, chess playing, engineering design, musical composition, and pattern recognition—are under way at other research centers. At least one computer now designs small standard electric motors (from customer specifications to the final design) for a manufacturing concern, one plays a pretty fair game of checkers, and several others know the rudiments of chess. The ILLIAC, at the University of Illinois, composes music, using I believe, the counterpoint of Palestrina; and I am told by a competent judge that the resulting product is aesthetically interesting.

Let me summarize as concretely as possible my assessment of the present and future state of the art and theory of heuristic problem solving. As of the present—1957:

1. Digital computers can perform certain heuristic problem-solving tasks for which no algorithms are available.

2. In doing so, they use processes that are closely parallel to human problem-solving processes.

3. Within limits, these machines learn to improve their performance on the basis of experience (not merely by memorizing specific patterns of successful behavior, but by reprogramming themselves in ways that parallel at least some human learning procedures).

On the basis of these developments, and the speed with which research in this field is progressing, I am willing to make the following predictions, to be realized within the next ten years:

1. That within ten years a digital computer will be the world's chess champion, unless the rules bar it from competition.

2. That within ten years a digital computer will discover and prove an important new mathematical theorem.

3. That within ten years a digital computer will write music that will be accepted by critics as possessing considerable aesthetic value.

4. That within ten years most theories in psychology will take the form of computer programs, or of qualitative statements about the characteristics of computer programs.

It is not my aim to surprise or shock you—if indeed that were possible in an age of nuclear fission and prospective interplanetary travel. But the simplest way I can summarize the situation is to say that there are now in the world machines that think, that learn, and that create. Moreover, their ability to do these things is going to increase rapidly until—in a visible future—the range of problems they can handle will be coextensive with the range to which the human mind has been applied.

What are the implications of this development? They are of at least three rather distinct kinds:

1. There will be more and more applications of machines to take the place of humans in solving ill-structured problems; just as machines are now being more and more used to solve well-structured problems.

2. There will be applications of machines to tackle ill-structured problems of such magnitude and difficulty that humans have not been able to solve them. (This is parallel to current applications of computers to the numerical solution of partial differential equations that lie beyond the capacity of hand methods.)

3. The research on heuristic problem solving will be applied to understanding the human mind. With the aid of heuristic programs, we will help man obey the ancient injunction: Know thyself. And knowing himself, he may learn to use advances of knowledge to benefit, rather than destroy, the human species.

In estimating the rates at which these developments will come about. it may be instructive to turn to a close analogy in the field of atomic energy. The implications of atomic energy are also threefold: (1) the generation of power to replace and augment power from conventional fuels; (2) the production of hitherto unrealizable concentrations of power (the primary peaceful application being thus far to the study of the interior of the atom); and (3) the use of radioactive materials as tracers for the study of physical and biological processes. The main point in drawing the analogy is that in both cases—computers and atomic energy—the usefulness of the first application hinges on economic calculations, while the significance of the other two rests mainly on their technical feasibility.

Atomic fuels will replace conventional fuels only when the capital costs per unit of energy-generating capacity are competitive with the capital costs of conventional plants. Computers for heuristic problem solving will replace executives only when the costs per unit of problem-solving capacity are competitive with the costs for executives. In neither case is it easy to make a forecast with available data, but it seems highly probable in both cases that the changeover, if it comes, will come gradually.

A substantial impact of heuristic problem solving on research (either in allowing us to tackle more difficult problems than humans now can, or in informing us how talented humans solve problems) is probably more imminent. Here—as in the parallel cases for atomic energy—the question will be very little 'How much will it cost?' and very much 'Can we do it?' It is neither a trivial nor a costless process to transfer from a productive scientist to a student the heuristic programs that make the former a powerful problem solver. To do this generally takes some twenty years of educational effort, and the undertaking is frequently unsuccessful. To reproduce in another computer a problem-solving program that has been learned and been proved effective by a first computer is a trivial matter. When machines will have minds, we can create copies of these minds as cheaply as we can now print books.

If what I have said still seems distant and speculative to you, I would like to recall to you again the precedent of Charles Babbage, who, always standing on the realities of the present saw the importance also of peering into the future and forecasting its shape.

> Perhaps to the sober eye of inductive philosophy, these anticipations of the future may appear too faintly connected with the history of the past. . . .
>
> Even now, the imprisoned winds which the earliest poet made the Grecian warrior bear for the protection of his fragile bark; or those which, in more modern times, the Lapland wizards sold to the deluded sailors;—these, the unreal creations of fancy or of fraud, called, at the command of science from their shadowy existence, obey a holier spell: and the unruly masters of the poet and the seer become the obedient slaves of civilized man.
>
> Nor have the wild imaginings of the satirist been quite unrivalled by the realities of after years: as if in mockery of the College of Laputa, light almost solar has been extracted from the refuse of fish; fire has been sifted by the lamp of Davy; and machinery has been taught arithmetic instead of poetry.[5]

PERHAPS this is an appropriate point to bring my speculations to a close and to summarize briefly the course of my argument. Up to the present time, operations research and the management sciences have been largely limited, by the nature of their tools, to dealing with well-structured problems that possess algorithmic means of solution. With recent developments in our understanding of heuristic processes and their simulation by digital computers, the way is open to deal scientifically with ill-structured problems—to make the computer coextensive with the human mind.

The energy revolution of the eighteenth and nineteenth centuries forced man to reconsider his role in a world in which his physical power and speed were outstripped by the power and speed of machines. The revolution in heuristic problem solving will force man to consider his role in a world in which his intellectual power and speed are outstripped by the intelligence of machines. Fortunately, the new revolution will at the same time give him a deeper understanding of the structure and workings of his own mind.

It is my personal hope that the latter development will outstrip the former—that man will learn where he wants to travel before he acquires the capability of leaving the planet.

REFERENCES

1. Charles Babbage, *On the Economy of Machinery and Manufacturers*, p. 45.
2. *Ibid.*, p. 193. [Quoted by Babbage from a *Note sur la publication, propose par le gouvernement Anglais des grands tables logarithmiques et trigonometriques de M. de Prony*, (1820)].
3. O. G. Selfridge, "Pattern Recognition and Modern Computers," and G. P. Dinneen, "Programming Pattern Recognition," both in *Proceedings of the 1955 Western Joint Computer Conference*, IRE.
4. "The Logic Theory Machine," *IRE Transactions* **IT-2,** 61–79 (September 1956); "Empirical Explorations of the Logic Theory Machine" and "Programming the Logic Theory Machine," *Proceedings of the 1957 Western Joint Computer Conference*, IRE; "The Elements of a Theory of Human Problem Solving," *Psych. Rev.*, in press.
5. Charles Babbage, *op. cit.*, p. 389.