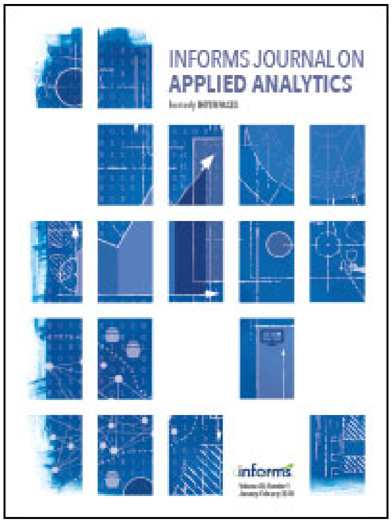

Interfaces

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

Two Heads Are Better than One: The Collaboration between AI and OR

Herbert A. Simon,

Two Heads Are Better than One: The Collaboration between AI and OR

HERBERT A. SIMON

Departments of Psychology and Computer Science
Carnegie-Mellon University
Pittsburgh, Pennsylvania 15213

The MS/OR profession has, in a single generation, grown from birth to a lively adulthood and is playing an important role in the management of our private and public institutions. This success should raise our aspirations. We should aspire to increase the impact of MS/OR by incorporating the AI kit of tools that can be applied to ill-structured, knowledge-rich, nonquantitative decision domains that characterize the work of top management and that characterize the great policy decisions that face our society.

During the past five years, we have seen a great deal in the public press and magazines about artificial intelligence and expert systems. The current widespread interest in expert systems may be partly attributed to the new general availability of minicomputers and personal computers that are capable of supporting them. The history of AI goes back much farther, however, almost to the beginnings of operations research. It is instructive to look at that early history in order to see why the two disciplines did not develop more nearly synchronously and with closer relation to each other.

The Origins of Operations Research

The contribution of operations analysis in World War II to solving a number of tactical problems relating to bombing raids and submarine detection suggested that the tools developed for these purposes, or related ones, might be equally useful if applied to the problems of managing civilian enterprises. We may take

0091-2102/87/1704/0008$01.25

PROFESSIONAL — COMMENTS ON ARTIFICIAL INTELLIGENCE

1945, when the war ended, as the birth year of OR, although TIMS and ORSA were not founded until 1952. The years from 1945 to 1955 saw not only the first civilian applications of the analytic tools of operations research, especially queueing theory, but also the invention of powerful new tools. Among these, by far the most important was linear programming, followed in a few years by integer programming, dynamic programming and control theory, critical path scheduling, and others.

Initially, the complexity of the problems that OR could handle was severely limited by the computational tools that were available: paper, pencils, and desk calculators. Many real-world problems for which the OR tools were otherwise appropriate were ruled out by the limits on computation. These limits were dramatically altered by the modern digital computer, which began to find a civilian market about 1950. With computers available, if the sky was not quite the limit, the ceiling was now very high; and it became higher each year as computers became larger and more powerful and as researchers in OR addressed the problem of improving the computational efficiency of the tools.

In addition to its invaluable contribution to computation, hence to the complexity of the problems that could be solved, the computer gave OR a new tool of analysis: simulation or modeling. Now, problems that were analytically intractable and not amenable to optimization could be approached by simulating system behavior numerically.

The symbiosis between OR and the computer has persisted down to the present time. While some new tools have been developed over the years, linear and integer programming and the other techniques introduced almost at the outset have continued to be the main weapons in the armory of OR. Improving the computational efficiency with which these tools could be applied has been the major focus of OR research.

The Origins of Artificial Intelligence

The history of AI is even more closely bound up with the introduction of the digital computer than is the history of OR. The computer was created as a number-crunching machine for solving very large arithmetic problems. If anyone attributed "intelligence" to it, it was only within the narrow limits of arithmetic computation, which it could obviously do very much faster and more accurately than we humans could.

The central insight that the computer was not limited to arithmetic calculation, but was a general-purpose symbol-processing system, derived from the stored-program concept. To interpret the instructions the computer found in its memory, it had to treat them not as numbers but as verbs in the imperative mode: "add!" or "multiply!" or "store!" From this followed the idea that the patterns stored in the computer could denote concepts of any kind and that the computer was capable of inputting and outputting such patterns, recording relations among them, storing them in associative structures, comparing them for identity or difference, and branching as the result of such comparisons.

By the year 1955, efforts had begun to

exploit the general symbol-manipulating powers of the computer for the performance of tasks that require intelligence when performed by people. The following year, 1956, is usually taken as the birth year of AI. In that year a computer program, the Logic Theorist, for the first time solved a problem (found the proof of a theorem) by the method of heuristic search.

Oddly, a good deal of the early history of AI took place in the basement of a business school. When the first computer, an IBM 650, was installed at Carnegie Institute of Technology, the electrical engineers did not want it, for they were afraid they would be asked to maintain it. The mathematicians did not want it, for they thought it beneath their dignity. Therefore, it found its home in the business school, and the first doctoral degrees awarded for dissertations in AI were degrees in industrial administration. The recipients included Allen Newell, one of the Founders of AI, Ed Feigenbaum, a pioneer in expert systems, and many others, all now prominent members of the AI and computer science communities.

In that setting, and in the decade after 1955, the tools of AI were applied side by side with OR tools to problems of management. Clarkson's dissertation [Clarkson 1962] was a description of an AI expert system (although the term had not yet been invented) that selected stock and bond portfolios for trust funds. Tonge [1961] constructed an AI program for balancing assembly lines. At the Westinghouse Electric Corporation, engineers constructed expert systems for the automatic design of motors, generators, and transformers.

Except for continuing research on heuristic methods for job-shop scheduling and engineering efforts like that at Westinghouse, where the inventors did not identify what they were doing as "artificial intelligence," these pioneering applications of AI methods to management problems were not followed up. After about 1960, AI and OR went their separate ways; whole new generations of scientists trained in each of these disciplines were largely unacquainted with the techniques provided by the other. Only when AI builders of expert systems began, about 1980, to invade the field that had traditionally been occupied by OR did each group begin to be aware again of the existence of the other.

> The symbiosis between OR and the computer has persisted.

Defining Some Terms

Before proceeding further, I will define what I mean by *operations research* and *artificial intelligence*.

Operations research may be defined as the application of optimization techniques to the solution of complex problems that can be expressed in real numbers. The criterion function, which determines what is to be optimized, must also be quantifiable. This definition is clearly too narrow to encompass all the things that operations research professionals do — for example, numerical modeling is excluded — but it characterizes the predominant emphasis upon formal mathematical models and optimization. It is therefore a useful definition for comparing and contrasting the approaches that operations research and artificial intelligence take to management problems.

By contrast, artificial intelligence is the application of methods of heuristic search to the solution of complex problems that (a) defy the mathematics of optimization, (b) contain nonquantifiable components, (c) involve large knowledge bases (including knowledge expressed in natural language), (d) incorporate the discovery and design of alternatives of choice, and (e) admit ill-specified goals and constraints.

This characterization of AI does not set very definite boundaries. It might be regarded more as a hunting license than as a proper definition. It emphasizes the aspiration of AI to deal with all the aspects of managerial decision making that stretch beyond the limits of classical OR. But it is preferable to stress continuities rather than differences. Nouns are the tyrants of the English language. When we introduce a new noun, we create the illusion that there must exist a recognizable entity corresponding to it, and we imply that to every pair of nouns there must correspond a pair of distinct entities.

Instead of differentiating between OR and AI, we need to confuse, blend, and synthesize them as much as possible. We need to build our professional institutions and organizations to use them together, supporting, reinforcing, and extending each other. But until we achieve this professional and organizational synthesis and create a management science whose scope is coterminous with the scope of management, my definitions will serve to indicate the rather different emphases that OR and AI have had and continue to have up to the present day.

Why Did AI Disappear from Management Science?

If, starting from their present status of regrettable isolation, AI and OR are to be brought into collaboration again, we need to understand the reasons why they strayed apart. There appears to have been no single cause, but a number of quite independent factors that cumulated to produce this estrangement.

One factor, although it may appear paradoxical, was the rapid increase in the power of computers and the simultaneous increase in the efficiency of OR computational algorithms. These developments made it possible for OR techniques to embrace an ever-widening range of practical management problems and thereby postponed the day when other tools, not limited to numerical optimization, would have to be added to the armory to encompass the whole range of problems of interest.

At the same time, the very limitations of computers during the early years of management science, and particularly the very small memories they had, put severe constraints on the expert systems that could be implemented. In the first decades of research on artificial intelligence, the main emphasis was placed upon gaining an understanding of the mechanisms of problem solving and intelligent inference. Attempts to use these mechanisms in expert systems that required large data bases waited until central processors with relatively large core memories began to be available.

The organization of the two disciplines also contributed to the separation of AI from OR. Most AI researchers were also active in other aspects of computer science. When computer science grew rapidly into an independent discipline, with its own professional organizations, journals, and university departments, AI became an integral component of the new discipline, especially at MIT, Carnegie-Mellon, and Stanford. AI researchers identified with the new professional and university organizations, loosening their ties with management science; and new recruits to AI had never had those ties.

The character of the research in the two fields also hastened the separation. OR was attractive to researchers with a mathematical bent, who could devote themselves to improving algorithms that had strong mathematical structure. AI research had a much more empirical and "engineering" cast, for it was usually impossible to prove strong theorems about AI systems. Their properties and capabilities had largely to be discovered by empirical testing.

> A good deal of the early history of AI took place in the basement of a business school.

During this period of growing separation, however, a variety of engineering groups in industry continued to develop programs to meet their practical needs, programs that we would now recognize as expert systems. Some of this work went on quite independent of the theoretical research in AI; in other cases, it was influenced by AI research. For example, chemical companies built systems for the automatic or interactive design of reaction paths for synthesizing organic chemicals. The idea for such systems was developed initially by the chemist E. B. Corey at Harvard [Corey and Wipke 1969], who derived their basic architecture from the General Problem Solver, an AI program designed around 1960.

What Needs to Be Done?

I have described a process of progressive isolation of artificial intelligence from operations research, and even from management science generally. Beginning about 1980, the eruption of a spate of expert systems, a number of them addressed to management problems, signalled AI's return to a concern with management [Harmon and King 1985]. The almost simultaneous appearance of microcomputers and personal computers greatly broadened public awareness of computers in general and the possibilities for artificial intelligence in particular. Articles about AI and expert systems began to appear frequently in newspapers and weekly and monthly magazines.

Expert systems and other applications of AI have been developed almost exclusively within the AI community and by a current generation of AI researchers who have little or no acquaintance with the classical OR tools. For example, branch-and-bound methods are an important component of integer programming in OR, while AI has long used a method of heuristic search called "best-first search." There is a close affinity between branch-and-bound and best-first search, but the connection between the two methods has only occasionally been noted.

The lack of knowledge of complementary OR techniques has not prevented the builders of expert systems from achieving a considerable measure of success in dealing with a number of real-world management problems. The success has been most marked in knowledge-rich domains, especially where qualitative as well as quantitative information has to be taken into account. For example, while OR has been able to provide a theoretical analysis of scheduling problems for very simple situations — a job shop with a very small number of machines and processes — automatic scheduling of large job shops with many machines and processes has been beyond the scope of exact optimization methods. A scheduling system, in such a situation, must be built upon a data base that contains a detailed description of machines, scheduling routes, and jobs. That kind of a description can be achieved with the help of the data-base languages that have been developed in artificial intelligence.

If artificial intelligence has now demonstrated convincingly its applicability to management science, this does not mean that it is best developed in isolation. On the contrary, the way is now open for fruitful collaboration between AI and OR. We can begin to think of expert systems and the data bases that support them as matrices within which we can apply an amalgamation of AI and OR analytic and problem-solving methods. Optimization methods can be applied to subproblems within larger environments that are too complex for them. Exact theoretical analyses of simplified problems can suggest procedures that will have good heuristic properties when applied to complex real-world situations.

Perhaps what is most important in making rapid progress towards our goals is that we adopt a problem-oriented point of view. We must let the problem that we are trying to solve determine the methods we apply to it, instead of letting our techniques determine what problems we are willing and able to tackle. That means we must be willing to apply our entire armory of techniques, both heuristic and algorithmic.

Adding AI methods to the kit of tools for management science opens up new domains of application. First, it provides powerful methods for making use of nonquantitative information. Building effective decision aids for management requires that we start, not with the particular data we have available in our information systems or with some particular techniques for manipulating those data, but with an analysis of the decisions that managers must make and of the kinds of data and analytic techniques that would be most useful in reaching those decisions. Much of the information that is most relevant to management decisions, especially at the higher levels of management, does not come in the form of numbers but in the form of sentences in natural language. AI deals with the computer as a general symbol-processing

After about 1960, AI and OR went their separate ways.

system that is as comfortable with information in natural language as it is with numbers.

There is a second direction in which AI should be especially useful in broadening the domain of management science applications. Most classical techniques of decision analysis are limited to making choices among a set of alternatives that are given, explicitly or implicitly, in advance. But a great deal of real-world decision making — perhaps most of it — is concerned with creating alternatives among which choices can be made. The activities that create new alternatives are usually called design activities, and although we most often apply the term *design* to the work of engineers and architects, it is equally central to the work of managers. Companies and the organizational structures within them must be designed. Investment alternatives must be discovered, that is, designed. Marketing strategies must be designed. Products and product lines must be designed. The methods of heuristic search used in AI are especially well adapted to the activity of design.

If management science is to bring together the tools of AI and OR so that they can be used in collaboration, then some obvious and important changes are needed in our university curricula. Students in business schools and in OR programs will need to gain more knowledge than they now do of the theory and methods of artificial intelligence. At the same time, students of computer science will need to learn more than they now do of the classical methods of operations research. And while we are broadening the curricula in these directions, we should take the opportunity to provide more problem-oriented instruction — instruction that starts with real problems and moves from them to appropriate tools — instead of relying too exclusively upon instruction that is tool-oriented.

Are there any steps we need to take in planning the future of TIMS/ORSA to accommodate this collaboration? One such step is to encourage the membership in our organizations of AI researchers who are concerned with management applications of their subject. But as the number of our AI colleagues increases, we must take care that we do not segregate them, or that they do not segregate themselves, in separate special-interest sections. Special-interest groups focused upon tool building have their role in such professional organizations, but their utility is proportional to the intensity of their interaction with their problem-oriented professional colleagues. Through such interaction we sustain the specialists' responsiveness to real-world needs and prevent their becoming wholly preoccupied with the honing of their already sharp tools.

The guiding principle is: mix them up! Meetings, and especially sessions, that bring AI and OR into the same room can provide a model for the future. We must make sure that we have all the ingredients in the TIMS/ORSA membership and then stir well.

Conclusion

The grey heads among us have seen TIMS/ORSA grow from an egg into a large and lively bird. We see the professions these organizations represent playing an important role in the management of our private and public institutions.

This success, achieved in about one human generation, should raise our aspirations. We should now aspire to extend the applicability of management science beyond the middle-management problems where it has had its main impact. To do that, management science must incorporate a still more general and powerful kit of tools that can be applied to the kinds of ill-structured, knowledge-rich, non-quantitative decision domains that characterize the work of top management and that characterize the great policy decisions on peace, energy, and environment that face our society.

Management science and operations research are a part of the great effort, often styled the Second Industrial Revolution, that is striving to understand and enhance intelligence. Joining hands with AI, management science and operations research can aspire to tackle every kind of problem-solving and decision-making task the human mind confronts.

References

Clarkson, G. P. E. 1962, *Portfolio Selection: A Simulation of Trust Investment*, Prentice-Hall, Englewood Cliffs, New Jersey.

Corey, E. J. and Wipke, W. T. 1969, "Computer-assisted design of complex organic synthesis," *Science*, Vol. 166, No. 3902 (October 10), pp. 178–192.

Harmon, P. and King, D. 1985, *Expert Systems: Artificial Intelligence in Business*, John Wiley and Sons, New York.

Tonge, F. M. 1961, *A Heuristic Program for Assembly Line Balancing*, Prentice-Hall, Englewood Cliffs, New Jersey.

Ed. note: This paper was originally delivered as the plenary address at the TIMS/ORSA Joint National Meeting, Miami, Florida, October 1986.