May 1969

THE RELEVANCE OF ROBOT RESEARCH TO ARTIFICIAL INTELLIGENCE

by
Bertram Raphael

Submitted for publication in the Proceedings of the Fourth Systems Symposium, Formal Systems and Non-Numerical Problem Solving by Computers, Case Western Reserve University, Cleveland, Ohio, November 19-20, 1968.

Artificial Intelligence Group
Technical Note 13

Project 7494

I INTRODUCTION

A major review of research in artificial intelligence recently stated that "The intensive effort being invested on the development of computer-controlled hand-eye and eye-cart devices is ... the most unexpected occurrence in A.I. research in the 1963-68 period."[1*] This paper takes the contrary view, which is that interest in the development of such robot-like devices is both a natural consequence of past developments, and a necessary stimulant to future research in the evolution of artificial intelligence and non-numerical problem solving.

By "robot" we do not refer, for example, to material-handling machines that mindlessly[†] step through preprogrammed sequences of actions. Instead, we consider a robot to be a computer-controlled mechanism that can interact with its real-world environment in an autonomous, reasonably intelligent[†] manner. A robot must be able to discover things about its environment by gathering information with its sensors, figure out the solutions to problems involving mechanical actions, and carry out these solutions by activating motors. Thus, a robot system must have sensory pattern-recognition abilities and general problem-solving capabilities. It must have well-designed logical and physical interfaces between the controlling computer and the external sensory and motor hardware devices. Perhaps most important, it must have a central computer program that functions as an executive system, maintains internal computer representations of the robot's environment, communicates

---

[*] References are listed at the end of this paper.

[†] Although some readers may object to anthropomorphic terms such as "mind" and "intelligent" applied to robots, we feel that such usage is convenient and we shall not apologize for it again in this paper.

with the human supervisor, and decides when to invoke the various sensory, problem solving, or motor functions.

We prefer not to include under the heading of "robots" remote manipulators that operate under the direct control of human operators.[2] However, at least one research group[3] views these man-machine systems as part of a continuum that includes fully autonomous robots at the other extreme. In this view, some of the techniques and conclusions for remote manipulators these scientists develop may well be relevant to future, more autonomous systems.

Another family of robot-related projects is largely concerned with modeling physiological processes. One project bases the decision-making functions of a simulated robot upon a decision tree model of ethological behavior in lower animals;[4] another determines major behavioral modes of activity by a simulation of neural activity in the reticular formation of a vertebrate;[5] and a third project proposes an "android" whose mental processes will be completely self-organized by means of adaptive networks of "probabilistic state variable" elements reminiscent of certain neural models.[6] Projects such as these may result in better understanding of certain physiological mechanisms. However, any such _a priori_ limitation to a particular technique cannot help but limit the effectiveness of the resulting problem-solving system.

The unconstrained heuristic programming approach, on the other hand, focuses on a _problem_ (rather than a _technique_) and tries to solve it by _any_ available techniques. The robot problem has been attacked in this direct manner by at least three major research groups: the Artificial Intelligence Groups at MIT, Stanford Research Institute, and Stanford

University. Descriptions of the specific emphasis and accomplishments of each of these efforts are available in the literature.[7-11] Since the intelligence built into such systems is limited only by the ability of the programmers, these heuristic robot systems have exhibited behavior more complex than that of any other autonomous robot systems in existence. In the next section we show that these projects have evolved from the main stream of artificial intelligence during the past decade.

II TRENDS IN ARTIFICIAL INTELLIGENCE RESEARCH

The history of artificial intelligence (AI) research is largely one of case studies; one picks a problem that appears to be just beyond current computer capabilities, and attempts to develop some approach that can bring about a solution. Problems have typically been chosen from the domains of formal mathematics, games and puzzles, or simple everyday reasoning problems. Domains for related (although somewhat separate) lines of development include simple fact retrieval and construction ("question-answering systems") and pictorial analysis ("visual pattern recognition").

As the AI research field matures, the effort devoted to certain classes of problems intensifies, and the new problems chosen increase in difficulty. In addition, two trends can be seen in the choice of research projects:

(1) A trend toward generality in approach, and the development of formal methods applicable to many problems rather than just one specific problem;

(2) A trend toward the choice of rich, complex problem domains, for which the integration of a variety of problem-solving

tools is necessary in order to achieve a satisfactory solution.

These trends place certain requirements upon the problems that must be chosen as test domains and milestones for progress. Robot research represents an extremely suitable problem area for these purposes, as we shall illustrate in Sec. IV below.

A. The Trend Toward Generality in Approach

The development of game-playing programs started with special-purpose machines for games like Tic-Tac-Toe and Nim, and eventually computer programs for Checkers and Chess. However, in parallel with developments in specific games, some general formal principles emerged that are applicable to many games--e.g., "minimax" techniques, the $\alpha - \beta$ heuristic, and the use of evaluation polynomials.

In automated mathematics, the AI research emphasis has been shifting from specific programs for plane geometry, trigonometry, and integral calculus, to the automation of proofs in predicate calculus, the logical system in which most of mathematics can be expressed. These proof techniques promise also to be applicable to several non-mathematical problem areas.

In the question-answering area, early programs dealt with specific subjects--e.g., questions about a set of baseball games. Later programs were concerned with representations of factual data from a variety of subject domains. Several current projects concentrate on word and concept associations in general, and on the combination of general inference mechanisms with broadly-based data structures.

Visual pattern-recognition research started with the application of simple models of neural networks to the recognition of hand-printed characters. This approach evolved into a general methodology: Preprocess the data to form vectors of significant features, then categorize these vectors by means of an adaptive classification algorithm. Recent work has established theoretical (as well as practical) limitations of certain algorithms of this type, and new heuristic scene-analysis techniques are being developed that will be more useful in a variety of problem situations.

Some workers in AI have attempted to develop models or theoretical principles that are independent of particular problem situations. This work includes studies of formal representational and problem-solving techniques, and of general graph-searching methods. One hopes that general theoretical principles can be developed that will have wide applicability in AI.

In all the above branches of AI--game playing, mathematical problem solving, question answering, pattern recognition, and especially in abstract theoretical studies--research scientists have a continuing need for specific problem domains in which to experiment with each newly developed idea. In choosing a test problem, the scientist must be careful not to fall into the trap of choosing a problem particularly well suited to the approach being tested (since his purpose is to test the generality of the approach). Another pitfall in choosing a specific problem situation is that one may be drawn into making <u>ad</u> <u>hoc</u> use of characteristics of the problem, thereby losing sight of the original goal of generality.

Thus AI research needs some problem situations that are sufficiently well defined that the natural way of stating a problem is independent of the problem-solving approach to be used. Also, such problem situations should be sufficiently rich that several different areas of AI research are relevant, respectively, to several different aspects of a problem. Robots provide one of the few known problem situations with these characteristics.

B. The Trend Toward Richness in Problem Situations

Early work in AI consisted almost exclusively of very specific, well-defined projects: a logic machine, a checkers program, a geometric analogy tester, etc. As these problems were solved, AI researchers started choosing richer and more complex problem domains. This shift in the nature of the problems is independent of the shift, discussed above, in the nature of the approaches taken in solving specific problems.

In the area of game playing, after seeing successful programs developed for moderately difficult games such as Checkers and Kalah, some AI workers have moved on to more difficult games such as Chess and, more recently, Go. Several others, however, have begun to develop "general" game-playing programs--programs that would be capable of solving any of a large class of games, given a suitable statement of the rules.

Symbolic mathematical programs now exist for algebra, trigonometry, geometry, and calculus. In addition to going after more difficult domains such as differential equations, some workers have begun to package a complete "mathematical laboratory"--a man/machine system that puts a large set of mathematical tools simultaneously at the fingers of a human mathematician. Some of the current work in theorem proving is similarly aimed toward systems with capabilities in many specific problem areas.

Problem solving in general is a meta-example of this trend toward richer problem situations. What task can possibly be richer than the problem of how to solve problems? This task, and the related question of how to structure information in general (file structures, the problem of representation, etc.), have been receiving increasing amounts of attention from the AI community.

The AI field has now progressed to a point where it can begin to consider practical, rather than toy, problem situations. There are several motivations for this step: to test the capabilities of existing AI techniques, to develop ways for combining several techniques in a single application area, to discover interesting new problems for AI to solve, and perhaps to make practical contributions to other fields of study. Thus AI techniques have been successfully applied to a major problem in organic chemistry, and are being studied in connection with the design of computer-operating systems.

We can now begin to characterize the kind of problem situation that would best satisfy the current needs of the AI research community. It should have several separate aspects--e.g., pattern recognition, problem solving, information representations, natural language processing--that each have significant continuing research interest. It should offer specific, well-defined test problems, such as level of performance on certain tasks, to use as measures of research progress. It should require generality, like a general game player or problem solver, so that efficient solutions would use the same programs or data structures for several diverse purposes in the system. And it should be sufficiently rich and open-ended to offer a tremendous range of problems of ever-increasing complexity. To a large extent, robot research possesses these characteristics.

III THE SEMANTICS OF ROBOT RESEARCH

In addition to the issues of generality and richness discussed above, robot research has one important advantage over virtually any other AI research domain: The absolute clarity of definition of the task environment. The fact that robot devices deal directly with the physical world places the roles of data representations, programming languages, computer configurations, etc., in their proper perspective.

Consider the ultimate semantic significance of the usual problem domains of AI research--e.g., abstract mathematical constructs, game situations, or the meanings of English sentences. Since the computer scientist cannot get a firm handle on such abstract concepts, his first job in studying them is usually to define representations, in terms of lists, arrays, or bit tables, for the class of entities of interest. He then specifies the input-output language by means of which his programs will communicate with the representations. Finally, he designs algorithms for manipulating the representations, and measures their performance by the effectiveness of these manipulations.

This approach often leaves one with some nagging uncertainties. Does the chosen representation correspond closely enough to the real problem? If not, then much of the rest of the work may be irrelevant. Does the complexity of the algorithms depend on the nature of the representation's input-output language? If so, then perhaps the solution is unfairly constrained by minor technical considerations such as available software systems. Finally, how strongly does the success of the system depend upon the particular representation? Perhaps most of the intelligence was required at the representational stage, thereby evading the main AI problem.

The ultimate semantic significance of a robot's problem domain is the physical world around the robot--and the robot operates directly upon this domain. The input-output language consists of control signals that interrogate physical sensors such as cameras and bumpers, and activate motors on wheels or arms.

Of course, a robot system may still make use of a computer representation of its environment. However, the sensors create a fundamental link between the representation and the environment it is supposed to represent. The performance of the system is not measured in terms of the effectiveness of algorithms upon representations; it is measured instead by the effectiveness of algorithms, possibly in association with representations, in solving problems in the <u>real</u> environment.

An important consequence of this intimate relationship between robots and the real world is the fact that the consequences of a robot action can never be predicted with absolute certainty. If a chess-playing program opens with the move P-K4, then it can <u>know</u> that the pawn is on King-4, and proceed with its play from there. If a robot tries to move three feet forward, on the other hand, its computer has a terrible time proving whether the movement was successfully completed. Perhaps the wheels slipped; a brick wall was in the way; someone slid the carpet back under the wheels; a motor gear was recently changed so that the scale factor is wrong; etc. In summary, a problem solver in a formal domain is essentially done when it has constructed a plan for a solution; nothing can go wrong. A robot in the real world, however, must consider the <u>execution</u> of the plan as a major part of every task. Unexpected occurrences are not unusual, so that the use of sensory feedback and corrective actions are crucial.

IV EXAMPLES

We shall illustrate the richness and potential significance to AI of robot problems by sketching a few examples. (The examples discussed below are designed to be suggestive of possible interesting problem domains, rather than specific proposed or on-going projects.)

A. Games

Games such as Chess provide interesting formal problems, but do not involve the meaningful sensory and motor interaction that is important in robot research. On the other hand, by casting a robot in the role of a piece that has some autonomy on the playing board, we can use the structure of a game situation to provide goals and benchmarks for significant robot studies.

Consider the problem of a spelunker trying to find his way out of a cave. His success is highly dependent upon his abilities to gather information and allocate resources. Thus his problem is qualitatively different from the formal problem of finding a shortest path, given a complete topological description of the cave. Similarly, a Chess game may look quite different from the point of view of, say, the Red Queen (as sketched by Louis Caroll) than from that of the opposing players. Several interesting robot problems can also be posed with much simpler games.

Consider, for example, the game of "fox and geese." (This game is played on a checkerboard. The geese are four checkermen on the first row, and the fox is a checker king on the last row. All other squares are empty. No captures are permitted. The fox tries to get across the board, and the geese try to block him.) Now, picture a robot "fox" in a large room whose floor is divided into 64 squares. (We might ask four programmers to take the roles of the geese--or perhaps we should have four more robots!)

The fox must play according to the rules--or, later, we may require him to learn the rules, and punish him each time he violates one. We may tell him what move the geese last made, or we can expect him to figure it out from what he can see. We can control the advantage in the game by modifying the shape of the board, or allowing an occasional extra move. Or, we can shift to a real-time competative situation by allowing the fox and geese to move continuously at prescribed maximum speeds (except that only one goose may move at a time). Can we design programs that will enable the robot fox to perform well, and learn to improve his performance, in such situations?

B. Pattern Recognition

Thus far, robot pattern-recognition research has been limited to simple three-dimensional scene analysis--the recognition of shapes of objects the robot must pick up or maneuver among. Once some of these problems are solved, we can begin to insert into the robot system some of the "standard" two-dimensional pattern-recognition abilities. Most such PR systems--e.g., character recognizers--require the patterns to be normalized in size, position, and orientation before they can be recognized. These conditions would have to be relaxed for a robot. After all, people in modern society would be severely handicapped if they could not recognize and read many signs--e.g.,  or  --from many distances and orientations.

Suppose we build a robot that can read. Think of the variety of fascinating projects that would become available: A hand-eye robot that plays Scrabble; a mobile robot messenger that can "Find Bert Raphael's office" by first reading his office number in the company telephone book; or a mobile robot that can go on a "treasure hunt," where each subtask leads to a sign containing the next instruction.

C. Computer "Understanding"[12]

One of the major open problems of AI is that of organizing and utilizing a large data base of common sense information about the world. The development of a program that can answer reading comprehension questions about children's stories has frequently been suggested as a next important step in AI. However, would it not be more significant if a mechanical system could draw upon its knowledge of the world in the course of solving meaningful physical problems, rather than only generating verbal responses? For example consider the task, "Find Bert Raphael's office," mentioned above. Suppose the robot reads

NAME	EXT	ALT	LOC
RAPHAEL, B.	4122	2312	K2072

in the office telephone directory. Considerable knowledge of abbreviations, the normal organization of directories, etc., is necessary to determine that K2072 is the office number. After making that determination, will the robot know that Raphael's office is probably on the 20th floor? After entering the elevator and discovering, by reading the buttons, that only three floors exist in the building, will it correctly choose the second? What is the nature of knowledge about office buildings, so that the robot will expect offices on the second floor to be arithmetically ordered, with even and odd numbers on opposite sides of the corridor? How can we design a robot that can ask for directions or advice only when it needs it, and then remember the general principles it learns for future use?

D. Tool Using

The use of tools has been described as a major characteristic that distinguishes man from the lower animals. Can we endow robots with tool-using capabilities? Let us consider briefly the nature of tools and their use.

One might define a tool, narrowly, as an implement with a particular function--e.g., a screwdriver is a tool for driving screws. Thus, we might build a robot with many appendages--a screwdriver, wrench, suction cup, ice pick, etc.--and require it to select the appropriate one for each specific task. Thus, tool selection would be reduced to a generalized pattern-recognition problem.

More generally, a tool is defined by its form rather than its function. For example, who has not used a screwdriver to open a paint can, or a pencil to help open an envelope? We are all familiar with widely useful "tools" such as paper clips, rubber bands, and chewing gum. How do we decide which tool to use when we need one?

Suppose a robot is taught, as part of its "understanding" of the world, some facts about the physics of simple machines such as the lever and the wedge. Can we program it to make use of such "knowledge" when it next has to move a heavy object? Is it unreasonable to expect a cleverly programmed robot to select as "tools" appropriately shaped objects that it finds lying about in its environment?

V CONCLUSIONS

This paper has described some of the motivation for the study and development of robot systems, from the point of view of AI research. It has explored how robots can be a useful experimental domain for results from several areas of AI, and has suggested some interesting AI tasks that are uniquely well suited for robot experiments. These justifications for robot research are in addition to the exciting rewards one may expect from the achievement of practical robot systems.

Of course robot research, like most areas of experimental science, has certain inherent difficulties. The construction of robots is an extremely expensive occupation, filled with frustrating and time-consuming engineering problems of little long-range interest. However, many of us who have been engaged in these activities for several years find the rewards, in terms of both past and projected accomplishments, well worth the investment.

VI ACKNOWLEDGMENT

The views expressed above were developed largely as a result of work supported at Stanford Research Institute under Contract No. F30602-69-C-0056 with the Advanced Research Projects Agency and the Rome Air Development Center. The author is indebted to several members of the staff of the SRI Information Sciences Laboratory for some of the ideas expressed here. In particular, the idea of a tool-using robot was suggested by Cordell Green, John Wensley proposed the robot "fox and geese" game, and John Munson suggested important distinctions between formal problem solving and robot research tasks.

REFERENCES

1. E. A. Feigenbaum, "Artificial Intelligence: Themes in the Second Decade," Proc. IFIP Congress 1968 (to be published).
2. E. G. Johnsen and W. R. Corliss, Teleoperators and Human Augmentation, Report SP-5047, NASA, Washington, D.C. (1967).
3. T. P. Sheridan and W. R. Ferrell, "Human Control of Remote Manipulators," Proc. IJCAI, pp. 483-494 (1969).
4. L. Friedman, "Robot Control Strategy," Proc. IJCAI, pp. 527-540 (1969).
5. L. L. Sutro and W. L. Kilmer, "Assembly of Computers to Command and Control a Robot," Proc. 1969 SJCC (to be published).
6. R. Lee, Wright-Patterson Air Force Base, Dayton, Ohio. Unpublished presentation, 1968 FJCC.
7. B. Raphael, "Programming a Robot," Proc. IFIP Congress 1968 (to be published).
8. J. McCarthy et al., "A Computer with Hands, Eyes and Ears," Proc. 1968 FJCC, Vol. 33, Pt. 1, pp. 329-338.
9. A Guzmán, "Decomposition of a Visual Scene into Three-Dimensional Bodies," Proc. 1968 FJCC, Vol. 33, Pt. 1, pp. 291-304.
10. J. A. Feldman et al., "The Stanford Hand-Eye Project," Proc. IJCAI, pp. 521-526a (1969).
11. N. J. Nilsson, "A Mobile Automaton: An Application of Artificial Intelligence Techniques," Proc. IJCAI, pp. 509-520 (1969).
12. M. L. Minsky, ed., Semantic Information Processing (MIT Press, 1969).