July 1970

AN EXPERIMENTAL MOBILE AUTOMATON

by

Charles A. Rosen

Paper to be presented to the American Nuclear Society Eighteenth Conference on Remote Systems Technology, Washington, D.C., November 1970 (proceedings to be published).

Artificial Intelligence Group
Technical Note 39
SRI Project 8259

This research was sponsored by the Advanced Research Projects Agency and the National Aeronautics and Space Administration under Contract NAS12-2221.

INTRODUCTION

For the past five years there has been underway at Stanford Research Institute a large program whose major objectives are:

(1) To investigate and develop techniques in artificial intelligence, and

(2) Apply these techniques to the control of a mobile vehicle enabling it to carry out, autonomously, incompletely specified tasks in a realistic laboratory environment.

We are thus developing an integrated system of computer programs and controlled hardware to serve as a test bed for research in machine intelligence.

The tasks selected were such that would normally require human intellect in sensing, problem solving, planning, and execution. By developing artificial intelligence techniques of sufficient generality, we hope to devise systems capable of replacing humans in situations that are either environmentally hostile, too remote for satisfactory communication and control, or that require very rapid and tireless response to sensed signals.

The specialized techniques in artificial intelligence that were selected for continued research were: natural-language communication with the system, question answering, visual scene analysis, problem solving and planning, and representations and/or models. Each of these techniques had to be integrated with the hardware, all subject to an overriding program which provided executive control.

As of the end of 1969 there was implemented a complete system, hardware and software, which could demonstrate performance of simple classes of tasks requiring autonomous behavior. Furthermore, this past work has lead to a much clearer view of the key problems and their relative difficulty of solution, necessary steps for useful applications in the future.

Other large programs aimed at the development of integrated intelligent systems are presently going on at MIT,[1*] Stanford University,[2,3] and the University of Edinburgh.[4-6] A series of reports[7-14] detail much of the SRI work; several published papers[15-17] present overviews.

---

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

DESCRIPTION OF THE SYSTEM

The Hardware System

The mobile vehicle (Figure 1), which is driven by two battery-powered electrical step motors, can move forward or backward, turn to the right or left. The robot is linked to the computer via two channels--a one-way UHF television channel, and a VHF telemetry channel for all other two-way flow of information. Motor-control information from the computer is received over the radio link, and is stored locally (on-board the vehicle) and routed to each of the motors. Optical sensors monitor the number of steps each motor makes, on-board circuitry counts these steps and reports back to the computer over the radio link the completion of each action. Other controlled electro-mechanical functions include the setting of brakes, and the control of several other step motors that provide pan and tilt movements of the head. Sensory equipment includes cat-whisker-actuated bump detectors, an optical range-finder, and a television camera. Picture information sensed by the television camera is sent back by radio link to the computer for processing.

The main function of the sensory equipment mounted on the vehicle is to provide environmental information that is required by the computer to build up models of the environment for subsequent use. A subsidiary function is to provide protection for the robot and other objects. Thus, when the robot bumps into something, the robot's tactile sensors actuate local circuitry to turn off drive motors and apply brakes. The computer is signalled when such a collision occurs and can override the stop order

if the plan it generates calls for this action. Future applications may require a number of such local actions to protect the machine and environmental objects using more sophisticated sensory equipment.

The optical range finder, based on the use of optical triangulation principles, has a range from 3 to 30 feet with approximately 10 percent accuracy. On command, the head can be turned in any given direction; a linear, top-to-bottom optical sweep produces a series of distance-measuring signals, which are sent back to the computer to be added to the stored model information. On command, also, a 240-line resolution, 16 gray level television picture can be sent back and temporarily stored in computer memory for subsequent analysis leading to knowledge of important objects and features of the robot's environment.

Crude navigation of the vehicle is accomplished by using the known number of steps that each wheel drive motor makes together with the knowledge of the initial position of the machine.

The work reported here has been done with an XDS-940 time-sharing computer fitted with 32K of core and a paging system using a large magnetic drum as swapping storage, thus effectively magnifying the size and number of programs that can be rapidly brought into operation. The computer was time-shared by a number of research workers simultaneously with the operation of the robot. A new computer system of considerably increased power and core storage, based on Digital Equipment Corporation PDP-10 and PDP-15 computers, is now being installed.

Artificial Intelligence Techniques

A relatively small number of artificial intelligence techniques have been selected for research, development and application. We believe that the selected techniques are sufficient for supplying a base on which one can build complex autonomous systems and which can later be expanded as the state of the art advances. Following is a brief review of each of the techniques.

Communication with the System. The researcher communicates with the computer driving the system by means of a teletypewriter. Although it is possible to use highly compact, specially coded messages that are unambiguously interpreted by a computer program, it would be highly desirable to use a natural language. We are now able[18] to issue commands to the system using a simple subset of English. A special computer program translates the experimenter's commands in English into statements in a secondary language (statements in the predicate calculus), wherein the normal ambiguity of ordinary English is eliminated. Information other than commands can also be entered into the system, automatically translated into the intermediate language, and made available to other functional programs. Additionally, information gathered and processed by the robot system can be retrieved by the experimenter on demand and typed out in simple English.

Question Answering. Many information-processing systems exist that permit many facts to be stored in a computer and retrieved by appropriate indexing schemes. It is highly desirable to be able to retrieve information that is not stored explicitly, that is, to be able to obtain information that is implied or can be logically deduced from the stored facts. For example, by storing the general fact that humans have hands, the system could derive the fact that Joe, Mary, John, etc. have hands, if the system also had the facts that Joe, Mary, John, etc. were the names of humans, without explicitly requiring the storage of the separate facts that each of them has hands. Clearly, the knowledge that these were the names of humans would be important facts to derive answers to many other related questions. A question-answering system, QA3,[19-21] has been programmed that incorporates a powerful logical-deduction mechanism, and using as facts, statements in predicate calculus derived from the experimenter's entries, and analysis of the robot's own sensory data.

Sensory Perception and Analysis. We have chosen just a few sensory devices from the host of such sensors that are currently available for obtaining information from the environment. The most important of these is visual perception and analysis. Ultimately we would like to take a series of pictures of the environment in which the robot lives, select and recognize important objects and the relations between them, and make this information available to other programs as required. At present (and perhaps in the long future), it is impossible to store the greatest part of the information that can be sensed. Therefore a selection process must be embedded in the method by which the primary information is analyzed. Thus in simple tasks, such as the robot planning a route to a specified point in the laboratory room, the visual perception and analysis system must determine the location of obstacles, the traversible regions and passageways, as well as the locations of walls, doorways, etc. In harder tasks the perception system must be able to identify various objects and to determine important properties, such as color, size, shape, relative position, etc. At present there are operational programs[22,23] that process the information taken in by the television camera, compute and delineate regions of constant intensity and then use these regions as descriptors to identify each major object in the scene. Figures 2a-f illustrate the sequence of operations.

Methods are still very crude but expansible. For example, we intend to add both color and depth information to increase the power of region analysis and thus to permit us to increase the number of different objects that can be so identified, especially in complex situations where some of the objects are partially hidden by others, or partially obscured by background detail. It should be noted that to do machine perception it appears quite necessary to have available to the visual perception program a great many facts about the particular environment in which the robot lives not obtainable by the visual sensor. As a simple example, it may be necessary for the system to be able to *deduce* that a dark rectangular shape discovered on a wall may be a doorway because measurements indicate that it has the correct "known" shape and dimensions for doorways, and is bounded by the "known" floor line.

Representations (Models). It has become increasingly clear that an automaton that must carry out incompletely specified tasks must have available to it a great number of facts about its environment. Equally important is the need for structuring available knowledge so that other computer functions, such as planning and problem solving, perception, question answering, etc., can have ready access to the body of knowledge without requiring selection and intervention by the human experimenter. The integrated structures of knowledge thus made available have been called representations and/or models. In our past work we have used

several types of such models. A geometric or grid model shown in Figure 3 divides up the space in a laboratory room into known empty areas, known locations of obstacles, and unknown areas. Another form of model, called a list model, associates important objects with various properties such as location, size, etc. These models are built up in computer memory using information obtained by the robot sensors, and in addition, entered by the experimenter via teletype. In a simple sense, the building up of these models containing carefully selected facts about the environment is a form of simple learning whereby planning and problem-solving functions can make use of the experience accumulated by the robot system itself.

Problem Solving and Planning. A major effort is being expended in developing programs[24,25] that can take an assigned task as input, break it down into subtasks and sub-subtasks, etc., until, at the lowest level, there is generated a sequence of effector (motor and switch) commands that are then transmitted and routed to the appropriate electromechanical devices on the robot. In essence, these programs attempt to solve the problem of attaining a specified goal, and in so doing construct a plan of action for the robot to follow. By executing, in sequence, the planned motor commands, the robot then attempts to carry out the plan. Ideally there should be considerable feedback by appropriate sensors to indicate that intermediate steps in the plan are, in fact, being carried out correctly, and corrective action would be planned and taken as necessary. In actuality, due to computational limitations, especially in picture processing, very little such feedback is now incorporated in the system. We intend to remedy this situation in future work.

Several methods of implementing problem solving and planning have been pursued. The most powerful method to date involves the use of a technique developed for automatic theorem proving, devised by J. A. Robinson.[26] Given a set of hypotheses, a computer program can prove truth or falsity of a given theorem in a constructive manner. This method has been adapted to robot problem solving by Green.[24] Unfortunately, its execution generates a large search tree which requires extensive computation. Considerable development of this method is continuing to find efficient methods of "pruning" search trees to reduce the necessary computation to an economic size.[27]

ROBOT TASKS

Programs have been written which permit the robot to do simple classes of tasks. A description of some tasks representative of several classes of such tasks, arranged in order of increased difficulty, follow.

Task One. The experimenter types in the x-y coordinates of the initial position of the robot and commands it in simple English, via teletypewriter, to go to a specified goal position. Initially no information about the environment is stored in the system's models. The robot turns in the computed direction of its goal, switches on its television camera and range finder, takes a picture and range information, which is transmitted back to the computer; programs analyze the data and determine the location of free space, corridors, and obstacles. It plans an efficient route to get to its goal, translates this plan into a sequence of motor commands, which it then attempts to execute. Should the robot bump into an obstacle it did not "see," it stops, backs up a little, faces the goal, takes another picture, and replans the rest of the journey. The new information regarding the location of the obstacle that was unknown in the first plan is entered into the system and used to update its models. The robot continues in this manner until it reaches its goal. After several trips in the laboratory, the system builds up within its models a reasonable description of its environment, which can be used for subsequent tasks. Thus, these simple tasks can serve as a means for exploring its environment.

Task Two. The experimenter asks the system to display the models that have been built up as a result of the execution of the previous tasks. A number of objects and their locations have been discovered and labeled. The experimenter then assigns a new task, which is to collect several selected objects together, not specifying how or where they are to be collected. A solution to the problem is automatically computed by the problem-solving programs, using the stored model information, together with axioms and routines, entered by the experimenter, which describe the logical consequences of pushing objects. A plan is developed in which the objects are collected by pushing several of them toward the location of one of them. This plan is translated into a series of motor commands and the robot attempts to carry these actions to a successful completion. In this series of tasks there is as yet no provision made for mistakes; that is, the successful completion of intermediate actions is not checked by taking new pictures and analyzing the state of affairs. In these instances failure to carry out the assigned task can and does occur.

Task Three. The robot is commanded to push a box off a shallow platform onto the floor. (Prior to this command there has been entered into the system some general axioms regarding the use of a ramp as a tool for permitting the robot to raise itself from floor height to the height of the platform.) The robot solves this problem using its problem-solving routines for generating a plan that involves locating the ramp, pushing the ramp over to the platform, riding up the ramp to get onto the platform, and finally pushing the box off the platform. Figures 4a-d illustrate the sequence of operations in this task.

Again, since there is no appropriate feedback, the robot quite often fails to do this task due to errors in exact location of the various objects, wheel slippage, or in imperfect execution of the

push routine. However crude the execution of this rather difficult task is, it does illustrate the use of a tool by an automaton, where the only information that has been entered into the system is quite general, that is, the principles by which this particular tool can be used. The problem-solving routines are sufficiently powerful to permit the system to apply these principles to specific situations not completely specified beforehand by the experimenter.

FUTURE WORK

We intend to continue intensive work to develop each of the above-mentioned artificial intelligence techniques. In particular, the visual perception research requires addition of more information gathered by the robot such as color, depth, and other known facts about the world, to become a really effective tool for building up useful models. New forms of representations of the data are being planned which can store in a unified manner necessary information in easily accessible forms. Problem-solving methods other than the formal theorem-proving techniques are being developed to augment and simplify the search procedures. Monitoring of actions will be added as a new function to check for attainment of intermediate goals, with provision to replan if necessary. Harder tasks are being devised, which will require the robot to expand its "world" and require the storing of many more environmental facts, as well as requiring increasingly more powerful problem-solving strategies.

POTENTIAL APPLICATIONS TO REMOTE SYSTEM TECHNOLOGY

Ideally a remote system would be able to do many classes of tasks autonomously. Thus the human operator would issue high-level commands such as "Explore the region surrounding the craft within a one-mile radius"; "Replace the dust filter"; "Build a bridge across crevice #5 able to support one ton"; etc. It would then be up to the automaton to carry out successively the necessary steps it would plan after it had gathered essential information and analyzed the specified task.

The present state of the art will not support the above capability. One can imagine, however, that such tasks could be performed by means of teleoperators under continuous control by humans. An evolutionary step upward would be for the human to break up the high-level tasks into a number of simpler tasks each of which could be performed autonomously. The human would monitor the successful completion of each subtask, perform the necessary analysis and reissue new subtasks if it were necessary to do so. Thus as the perception, problem solving and monitoring capabilities of automatons increase, more of these functions would be shifted from human to automaton, the division of labor being determined by technical and economic feasibility studies of the whole man-machine system.

SUMMARY

We have described an integrated system of computer programs that have been designed to control autonomously the activity of a mobile vehicle. Simple classes of tasks can be performed in a constrained laboratory environment, illustrating that it is possible for a computer-controlled machine to deal effectively with incompletely specified tasks that normally would require continuous human control and intervention. Although the tasks are relatively simple, there appears to be no insuperable difficulties in increasing the complexity to a point where practical utilization of the system is possible. An inordinate amount of computation is still required for accomplishing even simple tasks; thus, one main thrust of future research is to find simpler, more effective methods, especially for perception and problem solving.

ACKNOWLEDGEMENT

The research reported herein is the result of the efforts of the staff of the Artificial Intelligence Group and has been sponsored by the Advanced Research Projects Agency of the Department of Defense and monitored by the Rome Air Development Center under Contracts AF30(602)-4147 and F30602-69-C-0056, and by the National Aeronautics and Space Administration under Contract NAS 12-2221.

REFERENCES

1. Internal Memoranda, Project MAC, Massachusetts Institute of Technology, Cambridge, Massachusetts.
2. John McCarthy, Lester D. Earnest, D. Roj Reddy, and Pierre J. Vicens, "A Computer with Hands, Eyes, and Ears," Proceedings Fall Joint Computer Conference, Vol. 33, Pt. 1, pp. 329-338 (Thompson Book Co., Washington, D.C., 1968).
3. Jerome A. Feldman et al., "The Stanford Hand-Eye Project," Proceedings International Joint Conference on Artificial Intelligence, pp. 521-526a (Association for Computing Machinery, New York, May 1969).
4. R. J. Popplestone, "Freddy In Toyland," in Machine Intelligence 4, D. Michie and B. Meltzer, eds., pp. 455-462 (Edinburgh University Press, Scotland, 1969).
5. J. E. Doran, "Experiments with a Pleasure Seeking Automaton," Machine Intelligence 3, D. Michie and B. Meltzer, eds., pp. 195-216 (Edinburgh University Press, Scotland, 1968).
6. J. E. Doran, "Planning and Generalization in an Automaton/Environment System," Machine Intelligence 4, D. Michie and B. Meltzer, eds., pp. 433-453 (Edinburgh University Press, Scotland, 1969).
7. Charles A. Rosen and Nils J. Nilsson, eds., "Application of Intelligent Automata to Reconnaissance," First Interim Report, Contract AF30(603)-4147, SRI Project 5953, Stanford Research Institute, Menlo Park, California (November 1966), AD 817 189.

8. Charles A. Rosen and Nils J. Nilsson, eds., "Application of Intelligent Automata to Reconnaissance," Second Interim Report, Contract AF30(602)-4147, SRI Project 5953, Stanford Research Institute, Menlo Park, California (March 1967), AD 820 989.

9. Charles A. Rosen and Nils J. Nilsson, eds., "Application of Intelligent Automata to Reconnaissance," Third Interim Report, Contract AF30(603)-4147, SRI Project 5953, Stanford Research Institute, Menlo Park, California (December 1967), AD 827 938.

10. Nils J. Nilsson, Bertram Raphael, and Sven Wahlstrom, "Application of Intelligent Automata to Reconnaissance," Fourth Interim Report, Contract AF30(602)-4147, SRI Project 5953, Stanford Research Institute, Menlo Park, California (May 1968), AD 841 509.

11. Nils J. Nilsson, et al., "Application of Intelligent Automata to Reconnaissance," Final Report, Contract AF30(602)-4147, SRI Project 5953, Stanford Research Institute, Menlo Park, California (December 1968), AD 849 872.

12. Nils J. Nilsson, "Research on Intelligent Automata," First Interim Report, Contract F30602-69-C-0056, SRI Project 7494, Stanford Research Institute, Menlo Park, California (February 1969).

13. L. Stephen Coles et al., "Application of Intelligent Automata to Reconnaissance," Final Report, Contract F30602-69-C-0056, SRI Project 7494, Stanford Research Institute, Menlo Park, California (November 1969).

14. L. J. Chaitin et al., "Research and Applications--Artificial Intelligence," Interim Scientific Report, Contract NAS 12-2221, SRI Project 8259, Stanford Research Institute, Menlo Park, California (April 1970).

15. Nils J. Nilsson, "A Mobile Automaton: AN Application of Artificial Intelligence Techniques," Proceedings International Joint Conference on Artificial Intelligence, pp. 509-520 (Association for Computing Machinery, New York, May 1969).

16. Charles A. Rosen, "Machines That Act Intelligently," Science Journal, pp. 109-114 (October 1968).

17. Bertram Raphael, "Programming a Robot," IFIPS Congress Proceedings, Edinburgh, Scotland, August 1968 (to be published).

18. L. Stephen Coles, "Talking with a Robot in English," Proceedings International Joint Conference on Artificial Intelligence, pp. 587-596 (Association for Computing Machinery, 1969).

19. C. Cordell Green, "Theorem Proving by Resolution as a Basis for Question-Answering Systems," in Machine Intelligence 4, D. Michie and B. Meltzer, eds., pp. 183-205 (Edinburgh University Press, 1969).

20. L. Stephen Coles, "An On-Line Question-Answering System with Natural Language and Pictorial Input," Proceedings 23rd National ACM Conference, pp. 157-167 (Brandon/Systems Press, 1968).

21. C. Cordell Green and Bertram Raphael, "The Use of Theorem Proving Techniques in Question-Answering Systems," Proceedings 23rd National ACM Conference, pp. 169-181 (Brandon/Systems Press, 1968).

22. George E. Forsen, "Processing Visual Data with an Automaton Eye," in Pictorial Pattern Recognition, G. Cheng et al., eds., pp. 471-502 (Thompson Book Company, 1968).

23. Claude R. Brice and Claude L. Fennema, "Scene Analysis Using Regions," Artificial Intelligence (American Elsevier Publishing Company, Inc., New York, to be published).

24. C. Cordell Green, "Application of Theorem Proving to Problem Solving," Proceedings International Joint Conference on Artificial Intelligence, pp. 219-239 (Association for Computing Machinery, 1969).

25. Peter E. Hart, Nils J. Nilsson, and Bertram Raphael, "A Formal Basis for the Heuristic Determination of Minimum Cost Paths," IEEE Transactions on Systems Sciences and Cybernetics, Vol. SSC-4, No. 2, pp. 100-107 (July 1968).

26. J. A. Robinson, "A Machine-Oriented Logic Based on the Resolution Principle," Journal of the ACM, Vol. 12, No. 1, pp. 23-41 (January 1965).

27. Nils J. Nilsson, Problem-Solving Methods in Artificial Intelligence (McGraw-Hill Book Company, New York, to appear April 1971).