



STANFORD RESEARCH INSTITUTE
Menlo Park, California 94025 · U.S.A.

9 June 1969

Proposal for Research
SRI No. ESU 69-68

RESEARCH ON INTELLIGENT AUTOMATA

Part One--Technical Proposal

Prepared for:

Advanced Research Projects Agency
Washington, D.C. 20301

Prepared by:

C. A. Rosen
N. J. Nilsson
B. Raphael
R. O. Duda

Artificial Intelligence Group
Information Science Laboratory

Approved by:

David R. Brown

David R. Brown, Director
Information Science Laboratory

B. Cox for T. Meisling

Torben Meisling, Executive Director
Information Science & Engineering Division

Copy No. 16

Proposal for Research
SRI No. ESU 69-68

RESEARCH ON INTELLIGENT AUTOMATA

Part One--Technical Proposal

I BACKGROUND

Stanford Research Institute is now engaged in a project* to investigate and develop techniques in artificial intelligence and apply them to the control of a mobile automaton carrying out specified tasks in a realistic environment. These tasks would be such that normally require human intelligence in planning, sensing, and execution. Thus, an ultimate goal is to devise integrated 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 complex responses to sensed environmental signals. A second important goal is to develop artificial intelligence techniques that are sufficiently general to have wide applicability for industrial and government use. Numerous reports and papers^{1-13†} document the results of previous work on this project.

Our proposed work includes, specifically, research on the following major artificial intelligence topics: visual pattern recognition of everyday three-dimensional objects, problem solving, computer representations, and question answering. Each of these subjects is important in its own right, but also they are often intimately related and mutually supporting in implementing an integrated artificial intelligence system. Research on these topics, guided by the constraints and framework involved in developing such an integrated system, has a direction and purpose which we believe is significantly more productive than research done by individuals in relative isolation. The following outline will describe essential features of our proposed work.

II OUTLINE OF PROPOSED WORK

A. Problem Solving

In order for an automaton to perform a task efficiently, it should first plan or compute a sequence of actions that can be expected to complete the task. We call this process of computing actions problem solving. There are several different approaches to problem solving. We have

* Contract F30602-69-C-0056, Advanced Research Projects Agency, Department of Defense, and Rome Air Development Center.

† References are listed at the end of this proposal.

been following one based on formal logic; by finding a constructive "proof" to a theorem in a formal system that a certain task can be performed, we in fact find a sequence of actions that will perform it. Details on our previous work in theorem proving and its application to automaton systems can be found in several papers and reprints.⁹⁻¹⁵

During the next year we want to expand our work on the theorem-proving approach. Specifically we propose to:

- (1) Continue investigating the theory of efficient proof-finding methods in the predicate calculus. Several recent results for limiting the search for a proof show theoretical promise.^{16,17} We also want to continue recently begun work on how best to use proofs of known theorems in attempting to prove "analogous" ones.¹⁸
- (2) Test theorem-proving strategies by implementing them in a running theorem-proving program such as "QA-3."
- (3) Continue our investigation of the use of theorem-proving methods to construct programs of action for automaton problem solving.
- (4) Begin to experiment with ways to use theorem-proving methods in situations requiring "feedback" of new information acquired by sensors during task execution.

In addition, there are other approaches that might be useful in automaton problem solving. Three of these are:

- (1) The "production system" approach similar to that used by the General Problem Solver (GPS) of Newell, Shaw, and Simon.
- (2) The "procedural language" approach similar to that described by Fikes at Carnegie-Mellon University.¹⁹
- (3) A stored-action "decision tree" approach similar to that used by Friedman²⁰ to control a simulated robot, "Adroit."

Several questions that naturally arise are: What are the similarities and differences in these four approaches? Can they be combined? For what kinds of tasks and situations is each the most efficient?

We propose to begin studying some of these other methods and to implement in our automaton system any useful results.

B. Representations or Models

If the problem-solving system is to compute a sequence of actions having a desired effect in a real environment, it must have access to a stored representation of the environment, and it must have the ability to

compute the results of automaton actions upon the environment. We call the data structures that represent the environment a "model" of the world. An action can be simulated in a model, and its effect in the model can be taken to be an approximation of what the actual effect in the world would be.

Problem solving in our automaton system is accomplished by a hierarchy containing many levels of problem-solving programs. One important question is: Should each problem-solving program have its own model of that aspect of the "world" that it deals with or should a common model serve several? In either case another question is: In what form should model information be stored and how does this form depend on the level and method of problem solving? If several models are used, how should they interact? If it is appropriate to make a change in one model, should corresponding changes be made in others?

We have experimented so far with three types of models:

- (1) A "grid model" or map of the automaton's room telling which areas are full, which empty, and which unknown,
- (2) A "property-list model" telling the names and properties of various objects contained in the map, and
- (3) An "axiom model" that stores facts in the form of predicate calculus clauses.

The axiom model is appropriate for the high-level, deductive problem solving accomplished by theorem-proving methods; the grid model has been used by lower-level route-finding routines. Specifically, during the next year we wish to explore interactions between different types of models and to discover the advantages and limitations of each.

There are also problems concerned with the source of the information in models. Complete models can be provided by the programmer during runs through interactive teletype or light-pen input. Most importantly, we are interested in the case in which the sources of model information are the sensors of the automaton. Our current automaton system has visual, range-finding, and tactile sensors. Important questions demanding further research are: Which sensors should affect which models? How much and what type of sensory information should be stored in models? What sensory information should be actively sought rather than passively accepted? These questions begin to blend into the subject of Sensory Perception, to be discussed below.

C. Question Answering

An intelligent automaton--like an intelligent person--must have a large store of common-sense knowledge upon which to draw in order to solve any particular problem. "Question answering" is the study of the effective storage, retrieval, and use of such information.

An important part of this information is embodied in an awareness of the immediate environment. This "mental picture" is made up of the data structures we called a model in the previous section. However, many other kinds of data are also important for good problem-solving performance: memories of certain aspects of past situations, e.g., "I went to the movies yesterday"; descriptions of classes of objects that may not be part of the current environment, e.g., "Every ball is round"; certain laws of mathematics and physics, e.g., "If an object is moved from A to B it will no longer be located at A"; etc.

The main research issue in the question-answering area may be stated as follows: How can this wide variety of information be represented, stored, and indexed, so that the automaton can quickly identify the knowledge relevant to any particular problem-solving situation and then logically infer the consequences of that knowledge for the task at hand? Of course, this issue must be studied in close connection with the issues of how to represent the current state of the world ("modeling"), and how to perform logical inference ("theorem proving").

D. Sensory Perception

The primary source of sensory information in our automaton system is a vidicon camera. Since vision provides important information about the world in reasonably large quantities, it will continue to be predominant. Our present visual system provides information about full and empty areas in the room, but of course vision can potentially provide much more information. We propose to continue our investigation of the general problem of mechanizing visual perception, with emphasis on scene analysis techniques that exploit the mobility of the vehicle. We expect to extend our domain from the present well-controlled single-room environment to include at least corridors and other rooms.

Although it is something of an oversimplification, it is useful to distinguish between so-called low-level and high-level operations in scene analysis. Low-level operations are usually local and deal directly with the intensity values of the quantized picture. Typical low-level operators include spatial differentiators, edge detectors, texture operators, unguided edge followers, etc. High-level operations are distinguished by their employment of global criteria in an effort to analyze the scene and make sense out of the often imperfect information supplied by the low-level operators. Their search strategies embody the heart of any scene analysis program.

We propose to improve old and develop new low-level operators. Depth information from relative motion and stereoscopic pictures also appears to be promising, and we shall continue our work on area properties and region analysis.

In the area of high-level analysis, we have recently developed a decision-tree approach that promises to be an efficient means for analyzing simple scenes. In this approach we follow down a tree of possible

interpretations of a scene, keeping a record of confidences attached to each partial description. When the confidence of a picture drops too low, we "back up" the tree to a point where an alternative decision would have resulted in a higher-confidence description. Initial results here are promising, and we propose to investigate this approach further.

Because of the time required by the computer to process a complete scene, our present automaton system can only look at the world infrequently. Were the automaton to look at the world more frequently, fewer changes should occur from scene to scene, and it should be possible to take advantage of scene-to-scene correlations to reduce the time needed for processing per frame, and perhaps improve the reliability of performance as well.

We propose to investigate the general topic of the use of previous views in aiding scene analysis. This topic includes two problem areas of particular interest, search restriction and landmark identification. With many tasks, such as pushing objects or looking for particular objects, prior information is available to guide and restrict the visual search. Although a complete scene analysis could be performed, it should be possible to develop scene analysis techniques that are both more efficient and more reliable whenever such prior information is available. The landmark identification problem concerns the recognition of prominent, immovable objects seen previously, and the use of this information to reorient the automaton, correcting the accumulated dead-reckoning navigation errors. We have begun some simple investigations of these problem areas, and continuation of this work would be a major part of our vision research effort.

E. System Integration and Implementation

Although we have discussed problem solving, modeling, question answering, and perception separately, an important part of our project is the integration of these interacting processes into a single system. Thus progress in one of these areas often necessitates changes in the others. During the next year we expect to redesign major portions of the automaton software in order to consolidate many of the ideas that have arisen from our past work and to produce a new, more efficient, more flexible system. This redesign will take place in conjunction with the changeover to the proposed PDP-10 computer facility.

The computer changeover provides an idea opportunity for us to make major improvements in the implementation of the automaton system. These improvements will range from the choice of new programming languages to the redesign of mechanisms for communication among high-level problem-solving routines. These organizational improvements will complement the basic gains expected from the greatly increased memory and computing power of the new computer system. In addition, the use of a separate peripheral computer to monitor robot activity should result in long-term benefits in performance that more than justify the substantial implementation effort we anticipate investing in the new system.

Experimentation with the current version of the complete automaton system has stimulated our research in each of the four named areas. The environment of the current automaton system, although simple, still provides tasks that challenge our abilities. Simple route-finding tasks are easily accomplished, and we are beginning to experiment with tasks that involve pushing objects about and collecting them in specified places. Many of these will require the deductive abilities of our theorem-proving system. These kinds of automaton tasks will be continued next year, and we shall also initiate tasks requiring the automaton to use simple "tools."

Further complexity will be introduced by experimenting in the more complex environment provided by the corridors and offices adjacent to the currently used automaton laboratory area. These experiments will place much heavier demands on the versatility of the visual system.

As the environment and tasks are made more complex, we shall also study the feasibility of adding more sensory and motor capabilities to our present system. We are particularly interested in adding some speech-recognition capability (based on the results of speech research in other laboratories). Also it may be desirable to add a simple manipulator to increase the richness of the kinds of tasks that can be performed.

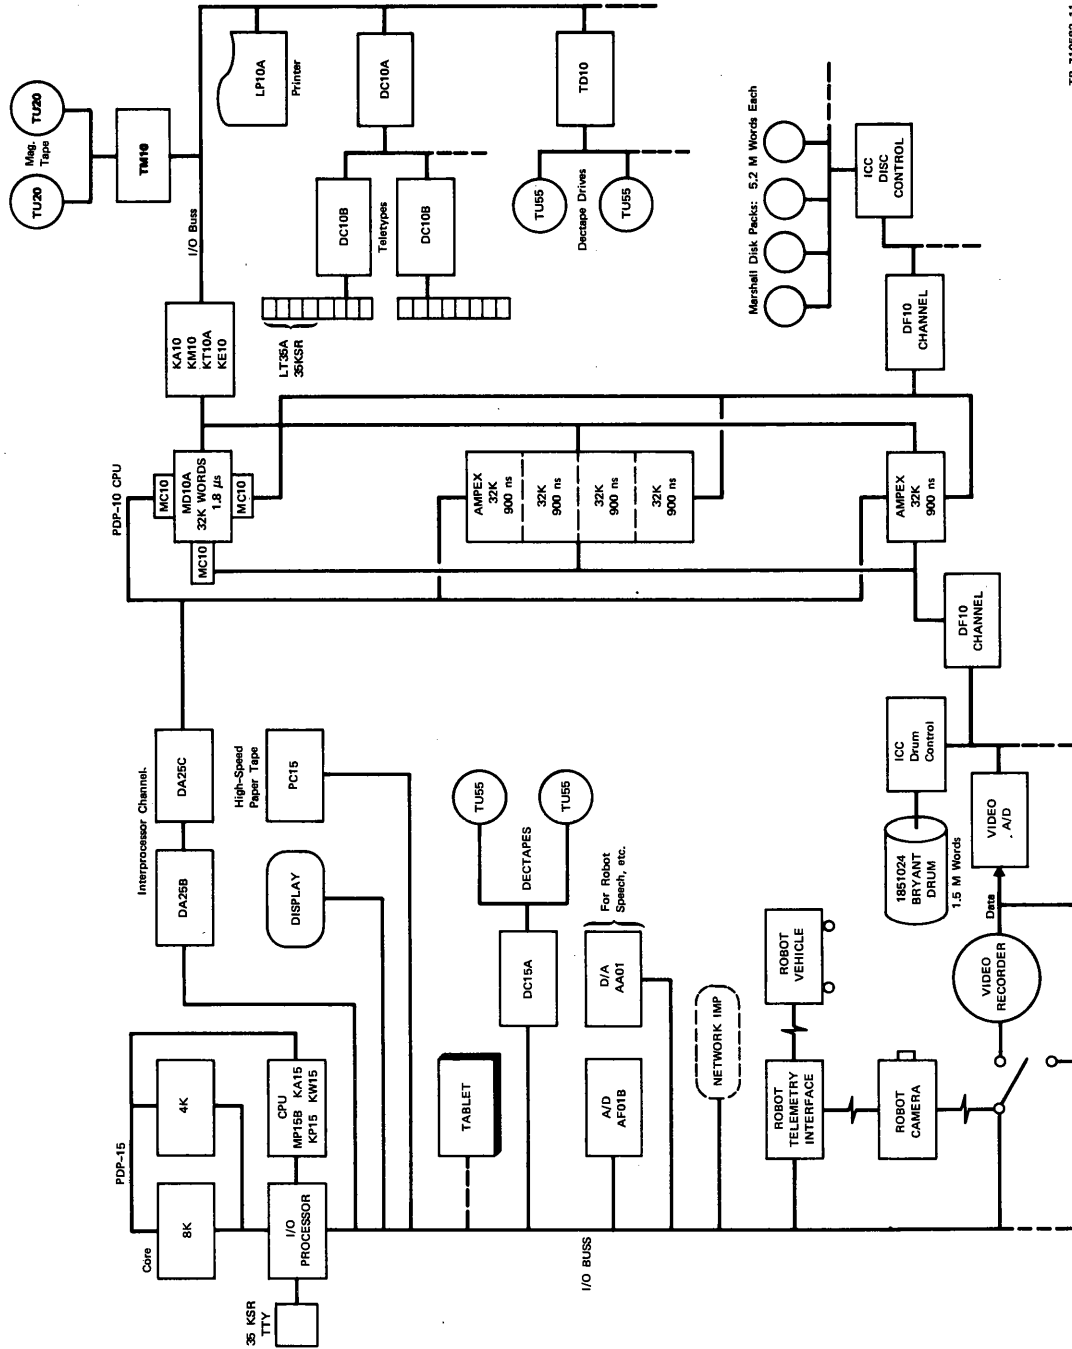
III EXPERIMENTAL FACILITIES

A. Computer

Until December 15, 1969, the Artificial Intelligence Group will continue to use an SDS 940 time-shared computer system. This system is leased and shared with numerous other projects. It has been evident for a number of months that this facility has serious limitations with respect to the size of direct-access core memory, and that significant response delays occur at peak use periods, especially when several large LISP programs are being used simultaneously. We are now negotiating the purchase of a new system based on the DEC PDP-10 computer, to be funded by ARPA, and dedicated to this project. We believe that the new system will clearly dominate other available systems suitable for our research purposes. We expect delivery of all the major components by the end of 1969 and to be operational shortly thereafter. Software changeover from the SDS 940 to the PDP-10 is being planned to start in early July 1969, and hopefully will be sufficiently advanced to permit operation with the new system when it has successfully passed its acceptance tests. We are making arrangements to "debug" software packages using one of several neighboring PDP-10 systems. The proposed configuration is shown in Fig. 1.

B. Special Equipment

We intend to make full use of the existing automaton vehicle together with its special equipment and interface with the computer. Relatively minor hardware changes are required for conversion of the interface with the new computer system. We are planning to make some minor modifications to the vehicle sensing and operating hardware. In particular, the



TB-710682-11

PDP-10 CONFIGURATION

most important additions will be to the visual system, for which we wish to add color-handling and other capabilities.

IV REPORTING SCHEDULE

An interim and final technical reports will be written, reporting on the results of the work performed.

V PERSONNEL

The following key personnel are expected to participate in the proposed research:

The following biographies were removed from this document prior to publication. All were part of the Information Science Laboratory of the Information Science and Engineering Division.

LEONARD J. CHAITIN, SYSTEMS PROGRAMMER
L. STEPHEN COLES, RESEARCH MATHEMATICIAN
RICHARD O. DUDA, SENIOR RESEARCH ENGINEER
CLAUDE L. FENNEMA, RESEARCH MATHEMATICIAN
RICHARD E. FIKES, RESEARCH MATHEMATICIAN
PETER E. HART, RESEARCH ENGINEER
PETER A. JOHANSON, RESEARCH ENGINEER
JOHN H. MUNSON, SENIOR RESEARCH PHYSICIST
NILS J. NILSSON, SENIOR RESEARCH ENGINEER
BERTRAM RAPHAEL, SENIOR RESEARCH MATHEMATICIAN
CHARLES A. ROSEN, MANAGER
RICHARD J. WALDINGER, RESEARCH MATHEMATICIAN
ROBERT A. YATES, SYSTEMS PROGRAMMER

REFERENCES

1. C. A. Rosen and N. J. Nilsson, eds., "Application of Intelligent Automata to Reconnaissance," First Interim Report, Contract AF 30 (602)-4147, SRI Project 5953, Stanford Research Institute, Menlo Park, California (November 1966), RADC-TR-66-822.
2. C. A. Rosen and N. J. Nilsson, eds., "Application of Intelligent Automata to Reconnaissance," Second Interim Report, Contract AF 30 (602)-4147, SRI Project 5953, Stanford Research Institute, Menlo Park, California (March 1967), RADC-TR-67-171.
3. C. A. Rosen and N. J. Nilsson, eds., "Application of Intelligent Automata to Reconnaissance," Third Interim Report, Contract AF 30 (602)-4147, SRI Project 5953, Stanford Research Institute, Menlo Park, California (December 1967), RADC-TR-67-657.
4. N. Nilsson, B. Raphael, and S. Wahlstrom, "Application of Intelligent Automata to Reconnaissance," Fourth Interim Report, Contract AF 30 (602)-4147, SRI Project 5953, Stanford Research Institute, Menlo Park, California (May 1968).
5. N. Nilsson, et al., "Application of Intelligent Automata to Reconnaissance," Final Report, Contract AF 30(602)-4147, SRI Project 5953, Stanford Research Institute, Menlo Park, California (December 1968).
6. C. A. Rosen and N. J. Nilsson, "An Intelligent Automaton," IEEE International Convention Record, Part 9 (1967).
7. B. Raphael, "Programming a Robot," Proc. IFIP Congress 68, Edinburgh, Scotland (August 1968).
8. G. E. Forsen, "Processing Visual Data with an Automaton Eye," in Pictorial Pattern Recognition (Thompson Book Company, Washington, D.C., 1968).
9. C. Green, "Theorem-Proving by Resolution as a Basis for Question-Answering Systems," Machine Intelligence 4, Meltzer and Michie, eds. (Edinburgh University Press, Edinburgh, Scotland; to appear 1969).
10. C. Green and B. Raphael, "The Use of Theorem-Proving Techniques in Question-Answering Systems," Proc. 1968 ACM Conference, Las Vegas, Nevada (August 1968).
11. L. S. Coles, "Talking with a Robot in English," Proc. Int'l. Joint Conf. on Artificial Intelligence (1969).
12. C. Green, "Application of Theorem Proving to Problem Solving," Proc. Int'l. Joint Conf. on Artificial Intelligence (1969).

13. N. J. Nilsson, "A Mobile Automaton: An Application of Artificial Intelligence Techniques," Proc. Int'l. Joint Conf. on Artificial Intelligence (1969).
14. C. Green and B. Raphael, "Research on Intelligent Question-Answering Systems," Scientific Report No. 1, Contract AF 19(628)-5919, SRI Project 6001, Stanford Research Institute, Menlo Park, California (May 1967).
15. B. Raphael, "Research on Intelligent Question-Answering Systems," Final Report, Contract AF 19(628)-5919, SRI Project 6001, Stanford Research Institute, Menlo Park, California (May 1968).
16. D. Luckham, "Some Tree-Paring Strategies for Theorem Proving," in Machine Intelligence 3, Michie, ed. (Edinburgh University Press, Edinburgh, Scotland, 1968).
17. P. A. Andrews, "Resolution with Merging," Journal of the ACM, (July 1968).
18. R. Kling, Internal SRI memos on "Reasoning by Analogy" (1967-1968).
19. R. Fikes, "A Study in Heuristic Problem-Solving: Problems Stated as Procedure," Proc. Fourth Systems Symposium, Case Western Reserve University, November 1968 (to be published).
20. L. Friedman, "Instinctive Behavior and its Computer Synthesis," Behavioral Science, Vol. 12, pp. 85-108 (1967).

13. N. J. Nilsson, "A Mobile Automaton: An Application of Artificial Intelligence Techniques," Proc. Int'l. Joint Conf. on Artificial Intelligence (1969).
14. C. Green and B. Raphael, "Research on Intelligent Question-Answering Systems," Scientific Report No. 1, Contract AF 19(628)-5919, SRI Project 6001, Stanford Research Institute, Menlo Park, California (May 1967).
15. B. Raphael, "Research on Intelligent Question-Answering Systems," Final Report, Contract AF 19(628)-5919, SRI Project 6001, Stanford Research Institute, Menlo Park, California (May 1968).
16. D. Luckham, "Some Tree-Paring Strategies for Theorem Proving," in Machine Intelligence 3, Michie, ed. (Edinburgh University Press, Edinburgh, Scotland, 1968).
17. P. A. Andrews, "Resolution with Merging," Journal of the ACM, (July 1968).
18. R. Kling, Internal SRI memos on "Reasoning by Analogy" (1967-1968).
19. R. Fikes, "A Study in Heuristic Problem-Solving: Problems Stated as Procedure," Proc. Fourth Systems Symposium, Case Western Reserve University, November 1968 (to be published).
20. L. Friedman, "Instinctive Behavior and its Computer Synthesis," Behavioral Science, Vol. 12, pp. 85-108 (1967).