EXPERIENTIALLY GUIDED ROBOTS[1]

E. WILLIAM MERRIAM and JOSEPH D. BECKER
Bolt Beranek and Newman Inc.
50 Moulton St.
Cambridge, Mass.

SUMMARY. This paper argues that an experientially guided robot Is necessary to successfully explore far-away planets. Such a robot is characterized as having sense organs which receive sensory information from its environment and motor systems which allow it to interact with that environment. The sensori-motor information which it receives is organized into an experiential knowledge structure and this knowledge in turn is used to guide the robot's future actions.

A summary is presented of a problem solving system which is being used as a test bed for developing such a robot. The robot currently engages in the behaviors of visual tracking, focusing down, and looking around in a simulated Martian landscape.

Finally, some unsolved problems are outlined whose solutions are necessary before an experientially guided robot can be produced. These problems center around organizing the motivational and memory structure of the robot and understanding its high-level control mechanisms.

This paper discusses a project which is attempting to develop the "mind" of' a robot which will be capable of experiencing its environment, storing sensori-motor information, and then using its accumulated knowledge to guide its future actions. In Section I, we describe the sort of behavior that an experientially guided robot might exhibit, and we give some reasons why we might want such a thing. Then, in Section II, we discuss the current state of our project, and in Section III we indicate some important issuses that are as yet unresolved.

I. MOTIVATION FOR BUILDING AN EXPERIENTIALLY GUIDED ROBOT.

The exploration of far space is characterized by problems which are likely to be solved only by employing an intelligent machine. For instance, the round-trip time required to travel from Earth to a relatively close planet like mars is aproximately two years, assuming current technology. We are still unsure of our ability to physically and psychologically

---

[1] This work was supported by the National Aeronautics and Space Administration under Contract Nrs. NASW-2236 and NASW-2572.

support a person in space for such a length of time. on the other hand, a machine would pose no such problems.

A problem that current machines do have, however, is that it is necessary to give them detailed instructions before they can carry out even simple tasks. These instructions have to come from Earth-bound humans who are analyzing the situation to determine what to do next. Unfortunately, the time required for a radio signal to travel from Earth to Mars is on the order of 10 minutes, so that a command-execution-feedback sequence would require about 20 minutes to be performed. Clearly, tasks that required many steps to complete would take an incredibly long time, and would probably not even be undertaken. When we attempt to explore planets beyond Mars, the time delays will be even greater, and the step-by-step execution technique will be totally useless.

A possible solution to this problem is to have some parameterized “canned” prodedures preprogramnied into the machine. Thus, a single command can be given which will elicit a whole sequence of actions and feedback information. If the canned prodedures have some degree of flexibility and a little bit of “intelligence”, then a wide variety of actions can be performed. For example, rocks of many different sizes can be picked up, looked at, and put in any desired place. This is the approach the robot project at the Jet Propulsion Laboratory [3] is taking, and it is probably essential I for the successful exploration of Mars.

Unfortunately, when exploring very distant planets with time delays measured in hours, even this approach will fail because of the relatively large number of commands and feedback information that will have to be transmitted. Furthermore, the information that we can preprogram into the robot about these planets will be far less reliable than that which we have for the Moon or for Mars. Therefore, our roving robot is likely to run into many more unexpected situations than in the case of the nearer places.

For these reasons, it is desireable to develop a machine that will be able to function with the bare minimum of commands from Earth. Such a robot should be capable of exploring an unfamiliar environment <u>on its own</u>. That is, upon arriving in a new environment, it should look around, note interesting rocks, landmarks, etc. and should send information pertaining to them back to scientists on Earth. It should then begin exploring the terrain and setting up experiments to collect scientific data. if something unexpected happens - either because of its own actions, the actions of wind, gravity, earthquake, living creatures, etc. - it must be able to take whatever action necessary and report the event back to Earth. It should be able to do all of this without a single command from Earth. This is not to say that it must ignore commands from Earth, but merely that it must be able to function well in a hostile environment without any such commands. Such a robot will prove invaluable for exploring far-away planets as well as in other applications nearer to home such as undersea mining, fire fighting, and hazardous rescue operations.

Before a robot will be able to successfully perform these actions, it will have to have acquired a considerable amount of knowledge about the world. Additionally, it must continually absorb new information about its surroundings and use that information to make intelligent decisions. This is the problem we are attacking. That is, we are trying to discover ways in which a robot can use sensori-motor information to build a knowledge base about the world, and then to use that knowledge to guide its future actions.

II. A ROBOT COMPUTER PROBLEM SOLVING SYSTEM. To provide an environment within which to test our ideas, we have simulated a Martian landscape through which a simulated robot roams. Since our main interest is in developing the intelligence capabilities of the robot, our use of simulations instead of real hardware does not unduly restrict our research. In fact, the development of suitable hardware is another research topic of its own and would only detract from our present goal.

A. The Robot and its World

The world which surrounds our simulated robot is a planar area which is partially bordered by the edges of three mountains and includes within it two hills and four craters (one a double crater). The plane (to which the robot is restricted) is randomly sprinkled with 50 rocks and 50 “unidentified objects” (which are small objects which resemble each other). Additionally, a single “instrument package” is included so that we may experiment with the effect of having a unique object in the environment.

This simulated Martian world provides an environment which is very rich visually, since it contains many objects both large and small, occurring in a variety of shapes and clusters. This informational richness, rather than any particular detail of its design, is the improtant feature of the simulated world as far as our study of robot intelligence is concerned.

The robot itself is a car-like vehicle with a tactile sensor which signals when it bumps into a mountain, hill, or crater. Additionally, it has a highly developed visual system which allows the robot to see the objects around it. Further, the blocking of one object higher than another is simulated so that the robot must deal with the phenomenon of an object disappearing behind another and other objects appearing as the robot moves about the world.

B. Behavior of the Robot

We describe here the behaviors that are currently programmed into our robot. They are each primitive in the sense that they will later be used to build up more elaborate behavior.

1. Visual Tracking

The robot can keep its eye on, or track, an object while it is moving. The basic operation of tracking is to notice that the tracked object has moved out of the center of the visual field, and to reposition the eye accordingly. In general, this repositioning can itself be a feedback process, but eventually all feedback processes must be composed out of non-feedback (endoqenous) primitives; otherwise, there would be a logical regress, and nothing would ever get done. We chose to make the repositioning operation be an endogamous process by making sure that the eye is always repositioned just enough to bring the object back into the center of the visual field.

2. Focusing Down

Once the robot can, through tracking, keep an object centered in its visual field, the next problem is to focus down on the object so as to pick up more and more of its detail. The details so obtained are used to build up a stored representation of the object. When, later, the object is to be visually recognized, the focusing-down procedure is reapplied to it, and the newly seen features are matched against those already stored.

In performing the focusing-down operation, the robot attempts to maximize the match between what it is currently seeing (held in an intermediate sensory memory structure) and what it once saw in the past (held in a long-term memory). In order to mediate this matching process over a period of time (and hence over several refocusings of the eye), a temporary representation of the object is built up in a short-term memory.

3. Looking Around

In order to gather information about its environment, the robot must look around at the scenery that surrounds it. Basically, the looking-around process is a solution to the problem of focusing down, with only one eye, on several recently-discovered objects all at once.

The scan paths obtained in our simulation when running the looking-around process look very reasonable, and seem to have a qualitative similarity to actual scan-paths that have been recorded in eye-movernent studies of human subjects. It would be very difficult to state more rigorous criteria for knowing when this program was behaving “correctly’.

Ultimately, we expect that the performance of any process will be evaluated, and if necessary corrected, by a higher-level process that makes use of it (e.g., in this case, a process that is trying to recognize where the robot is).

In this section, we have briefly described the primitive behaviors of our robot. Through programming these behaviors, we have come to grips with several issues that are important in developing an experientially guiiled robot. A more detailed discussion of our robot simulation and its behavior may be found in [1] and [2].

III. IMPORTANT UNSOLVED PROBLEMS. The behavior we have programmed into the robot is interesting from both the scientific viewpoint of understanding experientially-oriented intelligence and from the practical viewpoint of building a robot to accomplish useful exploration tasks. Our investigations have lead us to develop several memory structures for storing experiential knowledge and several control strategies for making use of it. However, the main benefit of our research so far is a qreater understanding of the problems that still remain to be. solved before we can hope to produce an experientially guided robot. We present here some of the more important problems with the hope that others may become intrigued with them and join us in developing an experientially guided robot.

A. Motivation

What motivates the robot to do anything at all? Why should it move? Why look around? Why track an object that happens to pass in front of the robot unexpectedly?

It seems clear that some motivations should be built into the robot system: such as “exploration” and “maintaining fuel supply”. But beyond this, the situation is unclear. That is, should only the main goals be built in and the robot forced to learn those activities (which may be different in different situations) that allow it to attain those goals, or should these actions be built in, too? In either case, there remains the question of what kind of control mechanism determines what goal or goals should be active at any given time. We and others have arrived at specific solutions to handle specific cases, but as of yet, there is no general solution or even accepted techniques for dealing with these problems.

B. Memory

At any given instant, the robot should have some memory structure. As every sensory or motor event takes place, that structure is modified, not by simply adding to it, but by changing existing structural elements. This changed structure, which is different than it has ever been before, is then used to guide the robot’s future actions. The original event may have no representation in this new structure, or the structure mav have been modified such that representations of previous events have been changed or obliterated. To “retrieve” them is now impossible. The representations of the original events have been irreversably modified by the accumulated events that took place before they occurred and by those which have occurred since.

Whereas we are convinced that the organization outlined above is correct, we still have few ideas on how to create such a structure either automatically or manually. Part of the problem arises because it is not yet clear what constitutes sensori-motor experience, so we don’t quite know what sort of “knowledge” we are dealing with. Mainly, though, the problem stems from lack of understanding of the workings of the processes which must operate on the memory structure. These are processes such as generalization, differentiation, deduction, and induction which give the robot the ability to react appropriately in situations that are not identical with those encountered in the past. Indeed, we are coming to believe that these processes shouldn’t exist as separate entities, but instead should be integral parts of the memory accessing mechanisms.

C. High-Level Control

Our notions of how to control a robot are tightly bound to notions of the executive systems of current digital computers. Whereas this may be a good starting point, the control system required for a robot will have to cope with simultaneouslv-running processes instead of pieces of several processes running sequentially. In any case, we have only vague notions on this subject, but the following questions seem relevant to it:

How are events coordinated which take place at the save time, but in different sensory or motor systems?

When are activities terminated? Should they be explicitly stopped or do they just die out because some other activity has taken over? What does it mean for an ill-defined process (such as “looking around”) to be “done”?

What causes an interrupt? How does one very interesting event override a very strong desire to do something completely unrelated? What makes an event “interesting”?

How is time dealt with? How are various sub-actions coordinated into a unified motor skill? flow does the robot know which actions can be speeded up, and which must be performed at a set pace?

The fact that we don’t yet have solutions to all of the above problems is certainly hindering the development of an experientially guided robot. We are, however, making steady progress on understanding the problems and are gradually developing solutions for them.

Sometimes these solutions are obviously not theoretically elegant, but by programming them, on our simulated robot, we gain greater insight into what the proper solution should be. In the meantime, we have solutions which may be useful in a practicable robot system.

IV. SUMMARY. In this paper, we have argued that an experientially guided robot is necessary if we hope to successfully explore far-away planets. Such a robot is characterized as having sense organs which receive sensory information from its environment and motor systems which will allow it to interact with its environment. The sensori-notor information which it receives from the environment and from itself is to be organized into an experiential knowledge structure, and this knowledge in turn used to guide the robot's future actions.

Additionally, we have presented a summary of a robot problem solving system which is being used as a test bed for developing our ideas. The robot currently engages in the behaviors of tracking, focusing down, and looking around in a simulated Martian landscape.

Finally, we have outlined some unsolved problems whose solutions are necessary before an experientially guided robot can be produced. These problems center around organizing the motivational and memory structure of the robot and understanding the mechanisms needed for the high-level control system of the robot. We hope that by discussing these problems, others will become excited about them and will join us in hopes of constructing an experientially guided robot.

References.

[1] Becker, Joseph D.; Robot Computer Problem Solving System; BRN Report Number 2316; September, 1972

[2] Becker, Joseph D. and Merriam, E. William; Robot Computer Problem Solving System; BBN Report Number 2646; September, 1973

[3] Lewis, Richard A. and Bejczy, Antal K.; "Planning Considerations For A Roving Robot With Arm"; Proceedings of the Third International Conference on Artificial Intelligence; Page 308-316; August, 1973