

## A COGNITIVE SCIENCE APPROACH TO IMPROVING PLANNING

Barbara Hayes-Roth  
Rand Corporation  
Santa Monica, CA. 90406

### INTRODUCTION

Planning is the predetermination of an intended sequence of actions aimed at achieving a goal. We all engage in planning for a variety of goals, ranging from everyday goals like performing a set of errands to more consequential goals like making a career change. Whether or not we achieve our goals depends in part on the quality of our plans.

During the past few years, my colleagues and I have been studying the cognitive processes people use for planning. When we began this work, most of the earlier scientific research on planning had focused on the development of automatic planning systems (e.g., Fahlman, 1974; Fikes, 1977; Sacerdoti, 1974, 1975). Other research had examined the role of plans in human behavior (e.g., Ernst & Newell, 1969; Miller, Galanter, & Pribram, 1969). However, little was known about the psychology of planning per se--how to identify effective planners, what special skills or strategies effective planners use, and what task factors impact on planning effectiveness. Because our long-range goal was to develop computer aids for human planners, we felt that understanding these psychological issues was an important first step. Accordingly, we embarked on a program of research designed to elucidate the cognitive processes underlying planning and to develop a computer aid that exploits cognitive strengths and compensates for cognitive weaknesses.

Of course, there are many different kinds of planning, depending upon the number of planners involved, the planning environment, the type of goals under consideration, the action options, etc. For our research, we chose to focus on individual planning of multiple-task sequences in a spatial environment. This task domain is well-defined and manageable. At the same time, it is general enough to apply to a variety of specific real-world planning tasks (e.g., planning travel itineraries, planning delivery routes, planning factory inspections, planning tactical missions). For our research, we wanted an instantiation of this task that was both realistic and familiar to the people who would serve as subjects in our experiments. We chose the following errand-planning task:

Given: a list of desired errands  
a map of the local environment  
starting and ending times  
starting and ending locations  
temporal constraints  
contextual information  
Plan: which errands to accomplish  
how much time to allocate for each errand  
in what order to conduct the errands  
by what routes to travel between successive errands.

Our research program comprises the following tasks:

1. Develop a cognitive model of the planning process.
2. Conduct experimental investigations of the model.
3. Apply the model in studies of individual differences in planning skill.
4. Apply the model in studies of planners' deficiencies and their vulnerabilities to task factors.
5. Infer principles for improving planning performance.
6. Design a computer aid around the inferred principles.
7. Implement the computer aid.
8. Test the computer aid in real planning environments.

We have completed tasks 1-5 and have recently begun working on task 6. This paper summarizes our work to date.

### A COGNITIVE MODEL OF THE PLANNING PROCESS

Our model of the cognitive processes underlying planning behavior exploits the theoretical architecture of the Hearsay-II speech-understanding system (CMU Computer Science Research Group, 1977; Erman & Lesser, 1975; Lesser, Fennel, Erman, & Reddy, 1975; Lesser & Erman, 1977; Hayes-Roth & Lesser, 1977). It also incorporates principles developed in the research on automatic planning systems and on the role of plans in human behavior mentioned above. The model has three basic components: specialists, a "blackboard," and a control regime. Each of these is discussed briefly below. (See Hayes-Roth & Hayes-Roth, 1979, for a more detailed discussion of the model.)

#### Specialists

Specialists are the mental processes that generate decisions for incorporation into the plan in process. For example, one specialist might generate a decision to establish a particular goal for the plan. Another might generate a decision to take a particular action toward achieving that goal.

We operationalize specialists as condition-action rules, similar to the production rules of Newell and Simon (1972). The condition part of the rule describes the circumstances under which the specialist can make a contribution to the plan. Ordinarily, the condition requires that some other planning decision has already been made. When that condition is satisfied, we say that the specialist has been "invoked." The action part of the rule describes the decision the specialist can contribute to the plan if it is "executed."

We assume that a given individual possesses many planning specialists. Some of them are generic and can

make contributions to all planning problems. Other specialists are domain-specific and can make contributions only to planning problems in their particular domains. We also assume that the many specialists an individual brings to bear on a planning problem operate independently. They do not communicate or influence one another's behavior directly. However, they can communicate and influence one another's behavior indirectly, as discussed in the next section.

### The Blackboard

The blackboard is a structured internal data base in which executed specialists record their decisions. All specialists can also inspect the blackboard and respond differentially to the presence of different kinds of information. In this way, specialists indirectly communicate and influence one another's behavior.

The model partitions the blackboard into five conceptual "planes" that distinguish the different kinds of decisions we think planners make. The meta-plan plane contains decisions about how to approach the problem, what kinds of problem-solving strategies to use, what kinds of policies should guide plan development, and what kinds of criteria should be used to evaluate tentative plans. The plan abstraction plane contains decisions characterizing the kinds of actions that should be included in the plan. The knowledge base plane contains data, assumptions, and knowledge about the world that might be useful in instantiating plan abstraction decisions. The plan plane contains decisions about the plan itself. These decisions are typically instantiations of plan abstraction decisions, based on information in the knowledge base. Finally, the executive plane contains decisions about how to sequence the execution of invoked specialists during the planning process. These decisions determine the order in which decisions are generated on the other planes of the blackboard.

The model further partitions each plane of the blackboard into different "levels of abstraction." To illustrate, the plan plane has four different levels of abstraction. The outcomes level contains decisions about the goals of the plan. The design level contains decisions about the general spatial-temporal organization of the plan. The procedures level contains decisions about the actions planned within that spatial-temporal organization. The operations level contains decisions about the low-level operations necessary to carry out those actions. The other planes have similar levels of abstraction.

The blackboard structure outlined above serves two functions. First, it embodies our model of the psychological categories of planning. Thus, it distinguishes our model from other planning models and provides one basis for evaluating the model's

psychological validity. Second, the blackboard structure improves computational efficiency by permitting specialists to restrict their inspection activities to those areas of the blackboard likely to contain information of interest.

### Control Regime

According to the model, planning proceeds in a series of "cycles." On each cycle, many specialists may be invoked. One specialist is scheduled to execute its action next. It does so, recording its decision at an appropriate location on the blackboard. The recording of a new decision signals the beginning of the next cycle. This process repeats until the planner has developed a satisfactory plan.

The process of scheduling one of the invoked specialists on each cycle is another important feature of the model. Most previous conceptions of the planning process imposed upon it a strict, hierarchical control regime. High-level abstract decisions were made first and refined by later decisions at successively lower levels of abstraction. By contrast, our model assumes an opportunistic control regime. Specialists are scheduled and decisions generated in highly variable orders determined by competing scheduling heuristics. We have concentrated on two scheduling heuristics, focus of attention and recency. The focus of attention heuristic recommends scheduling specialists that record decisions in pre-selected areas of the blackboard. The recency heuristic recommends scheduling specialists that have been invoked recently, for example, on the last one or two cycles. (Hayes-Roth & Lesser (1977) have recommended other heuristics, such as efficiency and efficacy.) We implement these heuristics by means of specialists that record relevant decisions on the executive plane of the blackboard.

The interaction of the focus of attention and recency heuristics can manifest a variety of specific control strategies, including the hierarchical strategy mentioned above. We believe that the flexibility embodied in the opportunistic control regime is both a more accurate model of the variability we observe in human planning behavior and a more powerful approach to planning in general.

### EXPERIMENTAL INVESTIGATIONS OF THE OPPORTUNISTIC MODEL

We conducted two kinds of experimental investigations of the planning model: psychological experiments and computer simulation experiments.

The psychological experiments provided support for several of the basic assumptions of the model, including the following: (a) that people make planning

decisions in each of the conceptual categories of the blackboard; (b) that people formulate plans at the postulated levels of abstraction; (c) that people develop plans with both top-down and bottom-up decision sequences; (d) that people effect alternative control strategies for planning; and (e) that people opportunistically exploit the information and constraints available during planning. (These experiments are discussed in detail in several reports: Goldin & Hayes-Roth, 1980; Hayes-Roth, 1980; Hayes-Roth & Hayes-Roth, 1979; Hayes-Roth & Thorndyke, 1980.)

The computer simulation was a LISP implementation of the model with about fifty specialists. The simulation served two functions. First, it demonstrated the sufficiency of the model to account for human planning behavior. The simulation produced plans and planning protocols similar to those produced by human planners. It also exhibited their characteristic strategic flexibility. Second, the simulation allowed us to explore some of the computational properties of the model, providing more general insights into distributed computation and heuristic control regimes. (The computer simulation is discussed in more detail in Hayes-Roth & Hayes-Roth, 1979, and in Hayes-Roth, Hayes-Roth, Rosenschein, & Cammarata, 1979).

#### INDIVIDUAL DIFFERENCES IN PLANNING SKILL

The first application of the opportunistic model was to individual differences in planning skill--why are some planners more effective than others? The model suggests three areas in which effective planners might differ from ineffective planners: their generation of decisions in different areas of the planning blackboard, their flexibility in distributing attention among the different areas of the blackboard, and their repertoires of specialists.

We evaluated these hypotheses by examining the planning processes of several planners with different skill levels. We assessed a planner's skill level based on the quality of the plans he or she produced. The quality of a plan was a composite score reflecting several interacting dimensions (e.g., efficiency, constraint satisfaction, temporal realism). Planners who achieved high plan scores were designated effective planners; those who achieved low plan scores were designated ineffective planners. We then examined the planning process of effective versus ineffective planners as revealed in thinking aloud protocols. Basically, we asked subjects to verbalize their thoughts while they formulated plans. We then analyzed these protocols, classifying statements as representing particular planes of the blackboard or levels of abstraction. Finally, we examined the relationship between planning skill and the frequency with which planners made different kinds of decisions.

The results supported all three hypotheses advanced above. Effective planners generated decisions in all areas of the planning blackboard, whereas ineffective planners generated primarily plan and plan-abstraction decisions. Effective planners also generated decisions at different levels of abstraction, whereas ineffective planners generated primarily low-level decisions. Effective planners showed greater attentional flexibility than ineffective planners. They more frequently shifted their focus of attention among the different planes of the blackboard and among different temporal loci within the plan itself. Finally, effective planners exhibited many more specialists than ineffective planners and they seemed to exploit powerful specialists more actively. (These results are discussed in more detail in Goldin & Hayes-Roth, 1980.)

#### PLANNERS' DEFICIENCIES AND THEIR VULNERABILITIES TO TASK FACTORS

We next applied the model to an analysis of general deficiencies in human planning and to the deleterious effects of task factors. Using protocol analysis methods similar to those described above, we identified three task factors that impede effective planning: constraints, complexity, and time stress.

Planners seem able to accommodate one or two simple time constraints. However, as the number of time constraints in a problem increases, planning effectiveness deteriorates. Time-constrained tasks, particularly those that appear late in the plan, rarely satisfy their constraints. The problem lies in planning strategy. The opportunistic model permits alternative planning strategies and, for heavily constrained problems, a constraint-based strategy is appropriate. Apparently, however, most of our subjects did not have a constraint-based strategy in their repertoires.

Planners also have difficulty coping with increasing problem complexity. As the number of tasks under consideration and the number of alternative possible plans increase, planners require disproportionately longer times to generate satisfactory plans. This problem also seems to lie in planning strategy. Instead of adopting a strategy which would restrict attention to a small number of the most promising alternatives, our subjects appeared willing to consider new alternatives throughout a planning session.

The third deficiency in human planning lies in the area of time estimation. Planners tend to underestimate the time required to execute planned actions and, as a consequence, to overestimate the number of actions they can execute in the available time. This tendency is exacerbated by time stress (the

time required to execute all the actions under consideration divided by the available time). Two factors contribute to this problem. Planners typically estimate time requirements at relatively high levels of abstraction. Because they fail to enumerate all time-consuming components of an action, they systematically underestimate the time required to execute it. Planners also respond to strong motivational factors. Their strong desire to accomplish all or most of the tasks under consideration biases their time estimates toward underestimation. (These results are discussed in more detail in Hayes-Roth, 1980.)

#### PRINCIPLES FOR IMPROVING PLANNING PERFORMANCE

Based on the opportunistic model and the empirical results summarized above, we developed the following principles for improving planning performance:

##### Criteria for Selecting Planners

1. Large-capacity working memory. The opportunistic model describes people's tendency to "jump around" the space of possible considerations while forming plans. This suggests that it may be important for planners to have large-capacity working memories in order to keep track of several aspects of a developing plan simultaneously.

2. Attentional flexibility. Our studies of individual differences in planning performance showed that good planners shift attention among different aspects of a planning problem more frequently than poor planners. Therefore, attentional flexibility may be another important characteristic of potential planners.

3. Strategic flexibility. Our studies of top-down versus bottom-up planning strategies showed the impact of particular planning strategies on the efficacy of the plans subjects produce and on the efficiency with which they produce them. In addition, some subjects appeared more willing than others to adopt alternative strategies. Therefore, strategic flexibility may be another important selection criterion.

##### What to Teach Planners

4. Concepts of abstract plans, meta-cognitive decisions, executive decisions, and knowledge-base decisions. Our studies of individual differences showed that good planners made decisions in all categories of the planning blackboard, whereas poor planners made only certain kinds of decisions. In particular, high-level abstract decisions, world knowledge decisions, metacognitive decisions, and executive decisions distinguished good planners from

poor planners. Therefore, planners should be taught the roles of these types of decisions in the planning process.

5. Domain-specific planning heuristics. Our studies of individual differences also showed that good planners had more different planning specialists than poor planners. Therefore, planners should be taught a variety of domain-specific planning heuristics.

6. Costs and benefits of opportunism. There is considerable evidence that most people employ some amount of opportunism in the planning process. Opportunism provides planners freedom to examine various aspects of a problem, investigate alternative plan configurations, etc. This enables them to discover solutions that more rigid approaches would obscure. On the other hand, opportunism requires additional time and may lead planners down unproductive, as well as productive, solution paths. Planners should be taught these advantages and disadvantages and how to exercise controlled opportunism.

7. General planning strategies. As mentioned above, different planning strategies are appropriate under different circumstances. Planners should be taught general planning strategies and the circumstances under which each is appropriate.

8. Judgment and time estimation. Most people show a strong tendency to underestimate the time required for planned actions. As a consequence, their plans are unrealistic and overrun the time available for execution. Planners should be taught cognitive methods for making such judgments more reliably and more accurately.

##### How to Train Planners

9. Provide explicit instruction. Explicit instruction appears to be a highly effective technique for training particular planning strategies and methods.

10. Induce illustrative experiences. Many planners seem to be able to generalize what they learn from one planning problem to subsequent, similar planning problems. Therefore, general strategies and methods can be taught by instructing planners how to use them on specific problems and providing opportunities for them to generalize them to similar problems.

11. Illustrate effective planning. Because our planning simulation effectively mimics the cognitive processes people use while planning, it may provide a useful model during training. The simulation could be programmed to illustrate planning strategies.

## Useful Aids for Planners

(See the following section.)

### ONGOING RESEARCH

As discussed in the introduction to this paper, we are particularly interested in developing computer aids to planning. We believe that, for the foreseeable future, people will play central roles in most important planning activities. Accordingly, an effective planning aid should exploit people's cognitive strengths and compensate for their cognitive weaknesses. We have recently begun working on the design of such an aid. Our work in this area focuses on a different instantiation of the same general planning task: project planning. We chose this task because it is an important real-world task that could benefit from the development of a planning aid and because we have contact with a variety of people who carry on project planning professionally. Our current design comprises the following components:

#### Goal Formulation

Our research suggests that planners suffer three deficiencies in the area of goal formulation. They do not formulate complete, well defined goals specifying all of the objectives, constraints, and policies the plan in progress should serve. They sometimes vacillate among conflicting goals, striving to satisfy different goals at different points in the planning process. They try to satisfy too many goals, given the available resources. The goal formulation component should help planners to articulate, prioritize, prune, and coordinate project goals.

#### Product Specification

Given a set of goals, the planner must generate functional specifications for project product(s). Presumably, development of these products would satisfy project goals. The problems in product specification are that planners may not generate complete specifications or they may not specify products that systematically satisfy project goals. The product specification component should help them to do so.

#### Task Analysis

Given a set of specifications, the planner must specify a set of project tasks whose execution would implement planned project products. Again, the problems are that planners may not generate a complete set of tasks or analyze them at a sufficiently low level of abstraction or that they may not specify tasks that systematically implement project products. The task analysis component should help them to do so.

## Resource Estimation

For each task under consideration, the planner must determine what resources are required to execute it. Our research suggests that planners are unduly optimistic about the number of tasks that can be accomplished with given fixed resources. The resources estimation component should help planners realistically assess the resource requirements of tasks under consideration.

## Resource Allocation

Given limited resources and alternative goals, the planner must determine how to allocate the available resources. Our research suggests that planners tend to spread resources too thinly across too many goals. The resource allocation component should help planners to formulate realistic allocation schemes and to perform cost/benefit analyses of alternative allocation schemes.

## Scheduling

The planner must schedule planned tasks in a way that provides adequate time for the execution of each task, insures completion of prerequisites by the time they are required, and minimizes slack time and other costs. The scheduling component should support these activities.

## Evaluation

Our research suggests that poor plan evaluation is a major impediment to effective planning. Poor planners do very little systematic evaluation of their tentative plans. Even good planners frequently decide arbitrarily among two or three final contenders. The evaluation component should at least assess whether the final plan meets criteria articulated by the planner. It should also assess the efficacy, efficiency and robustness of the plan in a simulated environment.

During the next few years, we plan to refine this design, implement it as a computer system, and evaluate its utility in the context of several Rand projects.

## REFERENCES

- CMU Computer Science Research Group. Summary of the CMU Five-year ARPA Effort in Speech Understanding Research, Technical Report, Carnegie-Mellon University, 1977.
- Erman, L.D. and V.R. Lesser. "A Multi-Level Organization for Problem Solving Using Many Diverse Cooperating Sources of Knowledge," Proceedings of the Fourth International Joint Conference on

- Artificial Intelligence, Tbilisi, USSR, 1975, 483-490.
- Ernst, G.W. and A. Newell. GPS: A Case Study in Generality and Problem Solving, New York: Academic Press, 1969.
- Fahlman, S.E. "A Planning System for Robot Construction Tasks," Artificial Intelligence, 1974, 5, 1-49.
- Fikes, R.E. "Knowledge Representation in Automatic Planning Systems," In A.K. Jones (ed.), Perspectives on Computer Science. New York: Academic Press, 1977.
- Goldin, Sarah and Barbara Hayes-Roth. Individual Differences in Planning Processes, N-1488-ONR, The Rand Corporation, June, 1980.
- Hayes-Roth, Barbara and Frederick Hayes-Roth. "A Cognitive Model of Planning," Cognitive Science, 1979, 3, 275-310.
- Hayes-Roth, Barbara and Frederick Hayes-Roth. Cognitive Processes in Planning, R-2366, The Rand Corporation, December, 1978.
- Hayes-Roth, Barbara and Perry W. Thorndyke. Decisionmaking During the Planning Process, N-1213-ONR, The Rand Corporation, October, 1980.
- Hayes-Roth, Barbara. Estimation of Time Requirements During Planning: Interactions Between Motivation and Cognition, N-1581-ONR, The Rand Corporation, November, 1980.
- Hayes-Roth, Barbara. Flexibility in Executive Strategies, N-1170-ONR, The Rand Corporation, September, 1980.
- Hayes-Roth, Barbara. Human Planning Processes, R-2670-ONR, Rand Corporation, December, 1980.
- Hayes-Roth, F. and V.R. Lesser. "Focus of Attention in the Hearsay-II Speech Understanding System," Proceedings of the Fifth International Joint Conference on Artificial Intelligence, Boston, Mass., 1977, 27-35.
- Lesser, V.R., R.D. Fennell, L.D. Erman, and D.R. Reddy. "Organization of the Hearsay-II Speech Understanding System," IEEE Transactions on Acoustics, Speech and Signal Processing, ASSP-23, 1975, 11-23.
- Miller, G.A., E. Galanter, and K.H. Pribram, Plans and the Structure of Behavior, Holt, Rinehart and Winston, Inc., 1960.
- Newell, A. and H.A. Simon. Human Problem Solving. Englewood Cliffs, N.J.: Prentice-Hall, 1972.
- Sacerdoti, E.D. "Planning in a Hierarchy of Abstraction Spaces," Artificial Intelligence, 1974, 5, 115-135.
- Sacerdoti, E.D. A Structure for Plans and Behavior. Technical Note 109, Stanford Research Institute, Menlo Park, California, August, 1975.