AD-A058 384 RAND CORP SANTA MONICA CALIF F/G 5/10
COGNITIVE PROCESSES IN PLANNING.(U)
AUG 78 B HAYES-ROTH, F HAYES-ROTH N00014-78-C-0039
UNCLASSIFIED RAND/WN-10268-ONR NL

1 OF 1
AD A058384

END
DATE FILMED
11-78
DDC

Contract No.: N00014-78-C-0039

LEVEL II

12

RAND/WN-10268-ONR

August 1978

COGNITIVE PROCESSES IN PLANNING

Barbara Hayes-Roth, Frederick Hayes-Roth

A WORKING NOTE
prepared for the
OFFICE OF NAVAL RESEARCH

This research was sponsored by the Personnel and Training Research Programs, Psychological Sciences Division, Office of Naval Research, under Contract No. N00014-78-C-0039, Contract Authority Identification Number, NR 157-411.

Rand
SANTA MONICA, CA. 90406

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE *(When Data Entered)*

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER Technical Report No. 1	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE *(and Subtitle)* (16) Cognitive Processes in Planning.		TYPE OF REPORT & PERIOD COVERED Technical Report 6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) (10) Barbara Hayes-Roth & Frederick Hayes-Roth		8. CONTRACT OR GRANT NUMBER(s) (15) N00014-78-C-0039
9. PERFORMING ORGANIZATION NAME AND ADDRESS The Rand Corporation 1700 Main Street Santa Monica, California 90406		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS NR 157-411
11. CONTROLLING OFFICE NAME AND ADDRESS Personnel and Training Research Programs Office of Naval Research (Code 458) Arlington, Virginia 22217		(11) 12. REPORT DATE August 1978 13. NUMBER OF PAGES 79
14. MONITORING AGENCY NAME & ADDRESS *(if different from Controlling Office)*		15. SECURITY CLASS. *(of this report)* Unclassified 15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT *(of this Report)*

Approved for public release; distribution unlimited.

(12) 93 p.

17. DISTRIBUTION STATEMENT *(of the abstract entered in Block 20, if different from Report)*

(14) RAND/WN-10268-ONR, TR-1

18. SUPPLEMENTARY NOTES

(9) Working notes,

19. KEY WORDS *(Continue on reverse side if necessary and identify by block number)*

Planning, cognitive model, multi-directional processing

20. ABSTRACT *(Continue on reverse side if necessary and identify by block number)*

(over)

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102 LF 014 6601

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE *(When Data Entered)*

296 600

Abstract

We propose a model of the planning process. Plannning is the predetermination of a course of action aimed at achieving a goal. The model assumes that planning comprises the activities of a variety of cognitive "specialists." Each specialist can suggest certain kinds of decisions for incorporation into the plan in progress. These include decisions about: (a) how to approach the planning problem; (b) what knowledge bears on the problem; (c) what kinds of actions to try to plan; (d) what specific actions to plan; and (e) how to allocate cognitive resources during planning. Within each of these categories, different specialists suggest decisions at different levels of abstraction. The activities of the various specialists are not coordinated in any systematic way. Instead, the specialists operate opportunistically, suggesting decisions whenever promising opportunities arise. We present a detailed account of the model and illustrate its assumptions with a "thinking aloud" protocol. We also contrast the model with earlier models of planning and discuss implications for future research.

PREFACE

This Working Note documents Rand research under Contract N00014-78-C-0039 with the Office of Naval Research. The objective was to construct and model a theory of human planning and control behavior.

ABSTRACT

We propose a model of the planning process. Planning is the predetermination of a course of action aimed at achieving a goal. The model assumes that planning comprises the activities of a variety of cognitive "specialists." Each specialist can suggest certain kinds of decisions for incorporation into the plan in progress. These include decisions about: (a) how to approach the planning problem; (b) what knowledge bears on the problem; (c) what kinds of actions to try to plan; (d) what specific actions to plan; and (e) how to allocate cognitive resources during planning. Within each of these categories, different specialists suggest decisions at different levels of abstraction. The activities of the various specialists are not coordinated in any systematic way. Instead, the specialists operate opportunistically, suggesting decisions whenever promising opportunities arise. We present a detailed account of the model and illustrate its assumptions with a "thinking aloud" protocol. We also contrast the model with earlier models of planning and discuss implications for future research.

Planning is a familiar cognitive activity. We all have many opportunities to decide how we will behave in future situations. For example, we plan how to get to work in the morning, where and with whom to eat lunch, and how to spend our evenings. We also make longer-term plans, such as what to do on our vacations, how to celebrate Christmas, and what career path to follow. Thus, planning influences many activities, from the most mundane to the most consequential, in everyday life.

We define planning as the predetermination of a course of action aimed at achieving a goal. It is the first stage of a two-stage problem-solving process. The second stage entails monitoring and guiding the execution of the plan to a successful conclusion. We refer to these two stages as planning and control. This paper focuses on the planning stage of the planning and control process. We have two main objectives: to characterize the planning process and to propose a theoretical account of it.

The planning process consists of a series of decisions regarding what to do and how to do it. Most decisions concern only limited aspects of the planned activities. They vary widely in the types of concerns they address and in the types of knowledge they exploit. Nonetheless, the decisions are not independent, but influence one another in important ways. An observation or decision regarding one aspect of the plan may influence a variety of subsequent decisions regarding other aspects of the plan. Thus, planning is largely an "opportunistic" process. The planner's initial observations

stimulate preliminary decisions. These, in turn, inspire subsequent observations and decisions, and the process repeats until an acceptable, comprehensive plan emerges.

The familiar task of designing an experiment illustrates our characterization of the planning process. The scientist must make a variety of decisions. He must define a goal, such as to test a specific hypothesis or to collect exploratory data. Keeping this goal in mind, he must choose or define independent and dependent variables, determine how many observations to make, create or select experimental materials, specify the experimental procedure, choose an appropriate laboratory apparatus, and adopt a model for statistical inference.

Obviously, many different kinds of knowledge influence these decisions. For example, the experimenter's knowledge of and hypotheses about the domain influence his choice of experimental goals. His long-term scientific goals, perhaps including a desire to understand an entire class of related phenomena, may also influence his choice of experimental goals. The scientist also presumably has task-specific procedural knowledge for designing experiments, collecting and analyzing data, drawing inferences from data, and so forth.

Although the scientist's decisions apparently concern different aspects of experimental design and reflect qualitatively different kinds of knowledge, they influence one another in important ways. For example, selecting a particular statistical design may subsequently constrain the scientist's

choice of sample size. Similarly, the current availability of a particular laboratory apparatus may influence the choice of independent and dependent variables. This may, in turn, influence the choice of the experimental goal.

In the following sections, we explore the planning process in more detail. In the first section, we characterize the planning problem and some of the complications that may arise during planning. In the second section, we characterize our theoretical approach and discuss an illustrative "thinking aloud" protocol of the planning process. In the third section, we present a model of planning. In the fourth section, we apply the proposed model to the thinking aloud protocol. In the fifth section, we compare the proposed model to previous models of planning. In the final section, we summarize our conclusions and discuss promising directions for future research.

THE PLANNING PROBLEM

Like other kinds of problem-solving (Cf., Newell & Simon, 1972), planning consists of a series of decisions made by a planner. The planner works with the following problem components: (a) an environment in which planning occurs; (b) a set of initial conditions, perhaps including some which suggest the need for a plan; (c) a set of possible actions which he can incorporate into a plan; and (d) a goal that he presumably can achieve by formulating and executing an appropriate plan. The planning problem, however, frequently entails a number of additional complications, as discussed below.

Problem components may not be fully specified for the executive. The planner may need to detect unsatisfactory or provocative initial conditions and, thus, infer the need for a plan. Similarly, the planner may need to discover alternative feasible actions and their likely effects. Finally, the planner may need to decide upon and define the goal. Designing an experiment illustrates planning with incompletely specified problem components. Scientists frequently examine recent empirical and theoretical findings to decide what issues merit further investigation. They also frequently develop new methods or paradigms for addressing the questions they formulate. Finally, scientists characteristically formulate both immediate and long-term goals for their research.

Problem components may be uncertain. It may not be possible for the planner to evaluate all environmental or initial conditions. It may not be possible to discover all possible actions or to determine the effects of particular actions. Preparing a study plan illustrates planning with uncertain problem components. Students rarely have advance knowledge of examination questions or acceptable responses. They also have limited knowledge of alternative study behaviors and their relative effects.

Problem components may change. Initial conditions, environmental conditions, possible actions, the effects of particular actions, and even goals may change. Further, changes in problem components may occur at any time: during planning, between planning and execution, or during execution of the plan.

Planning Social Security legislation illustrates planning with changeable problem components. While the original Social Security legislation embodied a reasonable programmatic approach to the goal of economic security, a number of unanticipated changes in the environment (e.g., dramatic increases in population and persistent inflation) undermined its effectiveness.

Because plans must be formulated in advance of execution, the planner does not have an opportunity to validate the presumed efficacy of particular components of the plan. Thus, the planner must formulate a complete plan (at some arbitrary level of detail) before he can evaluate its underlying assumptions or the efficacy of plan components. For example, a busy homemaker may assign dinner preparation to an adolescent child before leaving the house to perform necessary errands. Thus, the success of the homemaker's overall plan depends upon the child's success in preparing dinner (among other things).

Planning may involve extensive parallelism and coordination among intended actions. Only a limited class of problems can be solved by execution of a simple sequence of planned actions. Solutions to many problems require the coordination of multiple, simultaneous actions. Many problems also require that execution of particular actions be made contingent upon outcomes of other planned actions. Choreography illustrates the importance of parallelism and coordination in planning. When choreographing a dance for several dancers, the choreographer cannot simply plan movements for each of the

individual dancers. He frequently must plan the movements and gestures of several dancers to occur simultaneously or in particular variable sequences. Occasionally, he must coordinate their movements in more complex ways, as when a male dancer must lift or carry a ballerina.

Planning may have to be heuristic rather than algorithmic. It may not be possible to devise a plan whose execution guarantees achievement of the intended goal. Instead, the planner may need to devise a plan whose execution will probably achieve the goal or whose execution will approximate achievement of the goal. In such cases, the planner may need to formulate and comparatively evaluate several alternative plans. Strategic planning illustrates heuristic planning. Military conflict might erupt in the context of any of several widely differing scenarios. Further, the details of each such scenario are uncertain. Thus, rather than developing plans that guarantee deterrence of all potential threats, military strategists must develop plans that appear robust over some critical subset of the possible scenarios.

As characterized above, planning is considerably more complex than most of the cognitive functions studied by psychologists and computer scientists. Researchers have frequently ignored complex cognitive functions in favor of simpler, more tractable ones. Analyses of simpler functions presumably serve as "building blocks" in theories of more complex behavior. Although we respect the logic of this position, we follow Newell's (1973) suggestion that scientific progress on

complex cognitive processes requires "focusing a series of experimental and theoretical studies around a single complex task" (p. 303). Planning is well-suited to the approach suggested by Newell. In addition, recent theoretical advances in artificial intelligence (discussed below) make complex functions such as planning more amenable to investigation than they have been in the past.

THEORETICAL ORIENTATION

Our first assumption is that people plan opportunistically. That is, the planner does not take a systematic approach to formulating a plan. Instead, at each point in time, he works on whatever part of the plan appears most amenable to further development.

Two dimensions characterize the parts of a plan in progress: time and abstraction. The temporal dimension of a plan is obvious. A plan specifies a set of intended actions to be executed in some temporal configuration. In formulating a plan, the planner can work on sub-plans for initial actions, intermediate actions, or concluding actions. Note that a planner can treat the time dimension in absolute or relative terms. Thus, he can plan to perform a particular action at a specific time, or he can plan to perform it at an unspecified time relative to (e.g., before or after) some other action. The planner can also plan to perform actions at ill-defined times such as "at the end of the day."

The abstraction dimension distinguishes varying degrees of detail the planner includes in his consideration of the plan in progress. For example, he can formulate a very abstract plan, deciding only on a general approach to the problem at hand. At the other extreme, the planner can formulate a very detailed plan, deciding on all of the specific actions to be performed. (The abstraction dimension is ordinal or nominal, rather than continuous.)

We refer to the space of possible decisions bounded by time and abstraction dimensions as the "planning space."

The assumption that people plan opportunistically implies that the decisions they make can occur at non-adjacent points in the planning space. Further, a decision at one point in the planning space can influence subsequent decisions at other points in the planning space. Thus, a decision at a given level of abstraction specifying action to be taken at a given point in time may precede and influence decisions at either higher or lower levels of abstraction specifying actions to be taken at either earlier or later points in time.

The thinking aloud protocol in Figure 1 illustrates our characterization of the planning process. It was produced by a college graduate while planning a hypothetical day's errands. We have collected a total of thirty protocols from five different subjects performing six different versions of the errand-planning task. The protocol shown in Figure 1 is representative of this set. We chose it because it illustrates several of our points

nicely.

The subject began with the following problem description:

> You have just finished working out at the health club. It is 11:00 and you can plan the rest of your day as you like. However, you must pick up your car from the Maple Street parking garage by 5:30 and then head home. You'd also like to see a movie today, if possible. Show times at both movie theaters are 1:00, 3:00, and 5:00. Both movies are on your "must see" list, but go to whichever one most conveniently fits into your plan. Your other errands are as follows:
>
> > pick up medicine for your dog at the vet;
> > buy a fan belt for your refrigerator at the appliance store;
> > check out two of the three luxury apartments;
> > meet a friend for lunch at one of the restaurants;
> > buy a toy for your dog at the pet store;
> > pick up your watch at the watch repair;
> > special order a book at the bookstore;
> > buy fresh vegetables at the grocery;
> > buy a gardening magazine at the newsstand;
> > go to the florist to send flowers to a friend in the hospital

Note that the problem description specifies more errands than the subject could reasonably expect to accomplish in the time

available. The subject's task was to formulate a realistic plan indicating which errands he would do, when he would do them, and how he would travel among them.

Figure 2 shows the hypothetical town in which the subject planned his errands. Each of the pictures on the map symbolizes a particular store or other destination. The subject was quite familiar with both the symbology and the layout of the town. In addition, the map was available during planning.

We have numbered small sections of the protocol in Figure 1 to facilitate the discussion. Also for convenience, we refer to specific errands by the names of the associated stores or other destinations.

1 Let's go back down the errand list. Pick up medicine for the dog at veterinary supplies. That's definitely a primary, anything taking care of health. Fan belt for refrigerator. Definitely a primary because you need to keep the refrigerator. Checking out two out of three luxury apartments. It's got to be a secondary, another browser. Meet the friend at one of the restaurants for lunch. All right. Now. That's going to be able to be varied I hope. That's a primary though because it is an appointment, something you have to do. Buy a toy for the dog at the pet store. If you pass it, sure. If not, the dog can play with something else. Movie in one of the movie theaters. Better write that down, those movie times, 1, 3, or 5. Write that down on my sheet just to remember. And that's a primary because it's something I have to do. Pick up the watch at the watch repair. That's one of those borderline ones. Do you need your watch or not? Give it a primary. Special order a book at the bookstore.

2 We're having an awful lot of primaries in this one. It's going to be a busy day.

3 Fresh vegetables at the grocery. That's another primary. You need the food. Gardening magazine at the newsstand. Definitely secondary. All the many obligations of life.

4 Geez, can you believe all these primaries?

5 All right. We are now at the health club.

6 What is going to be the closest one?

7 The applicance store is a few blocks away. The medicine for the dog at the vet's office isn't too far away. Movie theaters--let's hold off on that for a little while. Pick up the watch. That's all the way across town. Special order a book at the bookstore.

8 Probably it would be best if we headed in a southeasterly direction. Start heading this way. I can see later on there are a million things I want to do in that part of town.

9 No we're not. We could end up with a movie just before we get the car. I had thought at first that I might head in a southeasterly direction because there's a grocery store, a watch repair, a movie theater all in that general area. Also a luxury apartment. However, near my parking lot

also is a movie, which would make it convenient to get out of the movie and go to the car. But I think we can still end up that way.

10 All right. Apparently the closest one to the health club is going to be the vet's shop. So I might as well get that out of the way. It's a primary and it's the closest. We'll start...

[The experimenter mentions that he has overlooked the nearby restaurant and flower shop]

11 Oh, how foolish of me. You're right. I can still do that and still head in the general direction.

12 But, then again, that puts a whole new light on things. We do have a bookstore. We do have...OK. Break up town into sections. We'll call it northwest and we'll call it southeast. See how many primaries are in that section. Down here we have, in the southeast section, we have the grocery store, the watch repair and the movie theater. In the northwest we have the grocery store, the bookstore, the flower shop, the vet's shop, and the restaurant.

13 And since we are leaving at 11:00, we might be able to get those chores done so that some time when I'm in the area, hit that restaurant. Let's try for that. Get as many of those out of the way as possible. We really could have a nice day here.

14 OK. First choose number one. At 11:00 we leave the health club. Easily, no doubt about it, we can be right across the street in 5 minutes to the flower shop. Here we go. Flower shop at 11:05. Let's give ourselves 10 minutes to browse through some bouquets and different floral arrangements. You know, you want to take care in sending the right type of flowers. That's something to deal with personal relationships.

15 At 11:10 we go north on Belmont Avenue to the Chestnut Street intersection with Belmont and on the northwest corner is a grocery.

16 Oh, real bad. Don't want to buy the groceries now because groceries rot. You're going to be taking them with you all day long. Going to have to put the groceries way towards the end.

17 And that could change it again. This is not one of my days. I have those every now and again.

(Figure 1)

Let's go with our original plan. Head to the southeast corner.

18 Still leaving the flower shop at 11:10. And we are going to go to the vet's shop next for medicine for the dog. We'll be there at 11:15, be out by 11:20. The vet's shop.

19 Proceeding down Oak Street. I think it would be, let's give ourselves a little short-cut.

20 Maybe we'll knock off a secondary task too.

21 Proceed down Oak Street to Belmont. Belmont south to the card and gift shop, or rather, to the department store. Cut through the department store to Johnson Street to the newsstand. Pick up our gardening magazine at the newsstand.

22 We're heading this way. We're going to make a definite southeast arrow.

23 Third item will be the newsstand since we are heading in that direction. Often I like to do that. I know buying a gardening magazine is hardly a primary thing to do, but since I'm heading that way, it's only going to take a second. Let's do it. Get it out of the way. Sometimes you'll find at the end of the day you've done all your primary stuff, but you still have all those little nuisance secondary items that you wish you would have gotten done. So, 11:20 we left the vet's office. We should arrive 11:25 at the newsstand. 11:30 we've left the newsstand.

24 Now let's start over here. We're going to be in trouble a little bit because of that appliance store hanging way up north. So we could; appliance store is a primary. It's got to be done.

25 All right, let's do this. This could work out. Market Square, we leave the Market Square exit of the newsstand up to Washington, arrive at the pet store, buy a toy for the dog at the pet store. We're there at 11:35, out at 11:40. Pretty good. 11:40. Proceeding east just slightly, up north Dunbar Street to the appliance store, we arrive there at 11:45, and we leave there, fan belt, leave at 11:50.

26 We're looking good. We've knocked off a couple of secondaries that really we hadn't planned on, but because of the locations of some stores they

are in the way that could be convenient.

27 Now it's 11:50, right near noontime.

28 And I think one of the next things to do, checking our primaries, what we have left to do, would be to go to the restaurant. And we can be at the restaurant at 5 minutes to noon. We're going to go down Dunbar Street, south on Dunbar Street to Washington east, to the restaurant which is located on the very eastern edge of the map. Meeting our friend there for lunch at 11:55, allowing a nice leisurely lunch. No, oh yeah. An hour, 12:55.

29 Now we've got to start being concerned about a few other things. We can pick up the car from the Maple Street garage by 5:30.

30 It's 12:55, done with lunch. Primary left to do, see a movie, pick up a watch, special order a book, and get fresh vegetables.

31 I would like to plan it so I can see the movie, pick up the vegetables, pick up my car, and then go home. Vegetables would rot.

32 So then with what we have left now to do is special order a book at a bookstore and pick up the watch at the watch repair.

33 So, I think we can make this a very nice trip. We're at the restaurant on Washington Avenue. Let's proceed west one block to Madison, south to Cedar Street. Cedar Street west right there at the intersection of Cedar and Madison is the watch repair. Pick up the watch at the watch repair. We should be at the watch repair by 1:05. Give us a good 10 minutes. 1:05 at the watch repair. Pick up a watch. We're out of there by 1:10.

34 Now I'm going to go just a slight back down Madison to one of the luxury apartments. I arrive at one of the luxury apartments at 1:15. I allow myself 15 minutes to browse. Two bathroom apartment. 1:30. Now I'm leaving that.

35 Next, I'm going to go west on Lakeshore, north on Dunbar, west on Cedar to the bookstore. And I will arrive at the bookstore at 1:35. Special order my book, 1:40.

36 From the bookstore I can go west on Cedar Street

just a hair, down Kingsway, to a second luxury apartment. Find out what's happening at that luxury apartment. And I'm there at 1:45, allowing myself another 15 minutes there, 2:00 we're out.

37 We've taken care of checking out 2 out of 3 luxury apartments. We ordered our book.

38 Now we do have a problem. It's 2:00 and all we have left to do is see a movie and get the vegetables. And that's where I think I've blown this plan. I've got an hour left there before the movie.

39 So the best way to eliminate as much time as possible since we are now located at the Cedar Lakeshore apartments. That's not going to be...

40 If I go get the groceries now, it's not really going to be consistent with the plans throughout the day because I've been holding off on the groceries for rotting. If I take them to a movie...Vegetables don't really perish like ice cream.

41 We leave the luxury apartment on Lakeshore, proceed due east to Dunbar, and we're at the grocery store at 2:05. 2:05 at the grocery store. Hunt around for fresh vegetables, and we can give ourselves 20 minutes there. So we leave there at 2:25.

42 We leave there and we proceed up Dunbar, north to Cedar, Cedar west to the movie theater.

43 We probably arrive at the movie theater at 2:35. 2:35 we arrive at the movie theater which still gives us 25 minutes to kill before the next showing. But that's that. We're going to have to simply do it. I'm going to have to go with it for right now.

44 The plan seems to have worked well enough up until then. We made better time than we had thought. That happens in life sometimes. How did I get here so fast?

45 2:25. We catch the 3:00 showing. We leave there at 5:00. Proceed immediately down Johnson, up Belmont to the parking structure, and we're there at 5:05 at the parking structure. We had to pick it up by 5:30.

46 Got everything done, the only problem being having

a little bit of time to kill in that one period.

47 You could have stretched out, to make things fair, you could have said, well, okay, I'll give myself an hour and 15 minutes at lunch, but as I did plan it, I did come up 30 minutes over. 25 minutes there. And that's a little bit of, when that happens you feel bad. You remember the old Ben Franklin saying about don't kill time because it's time that kills us. And I hate to have time to waste. I've got to have things work very nicely.

Figure 1. Thinking Aloud Protocol from the Errand-Planning Task

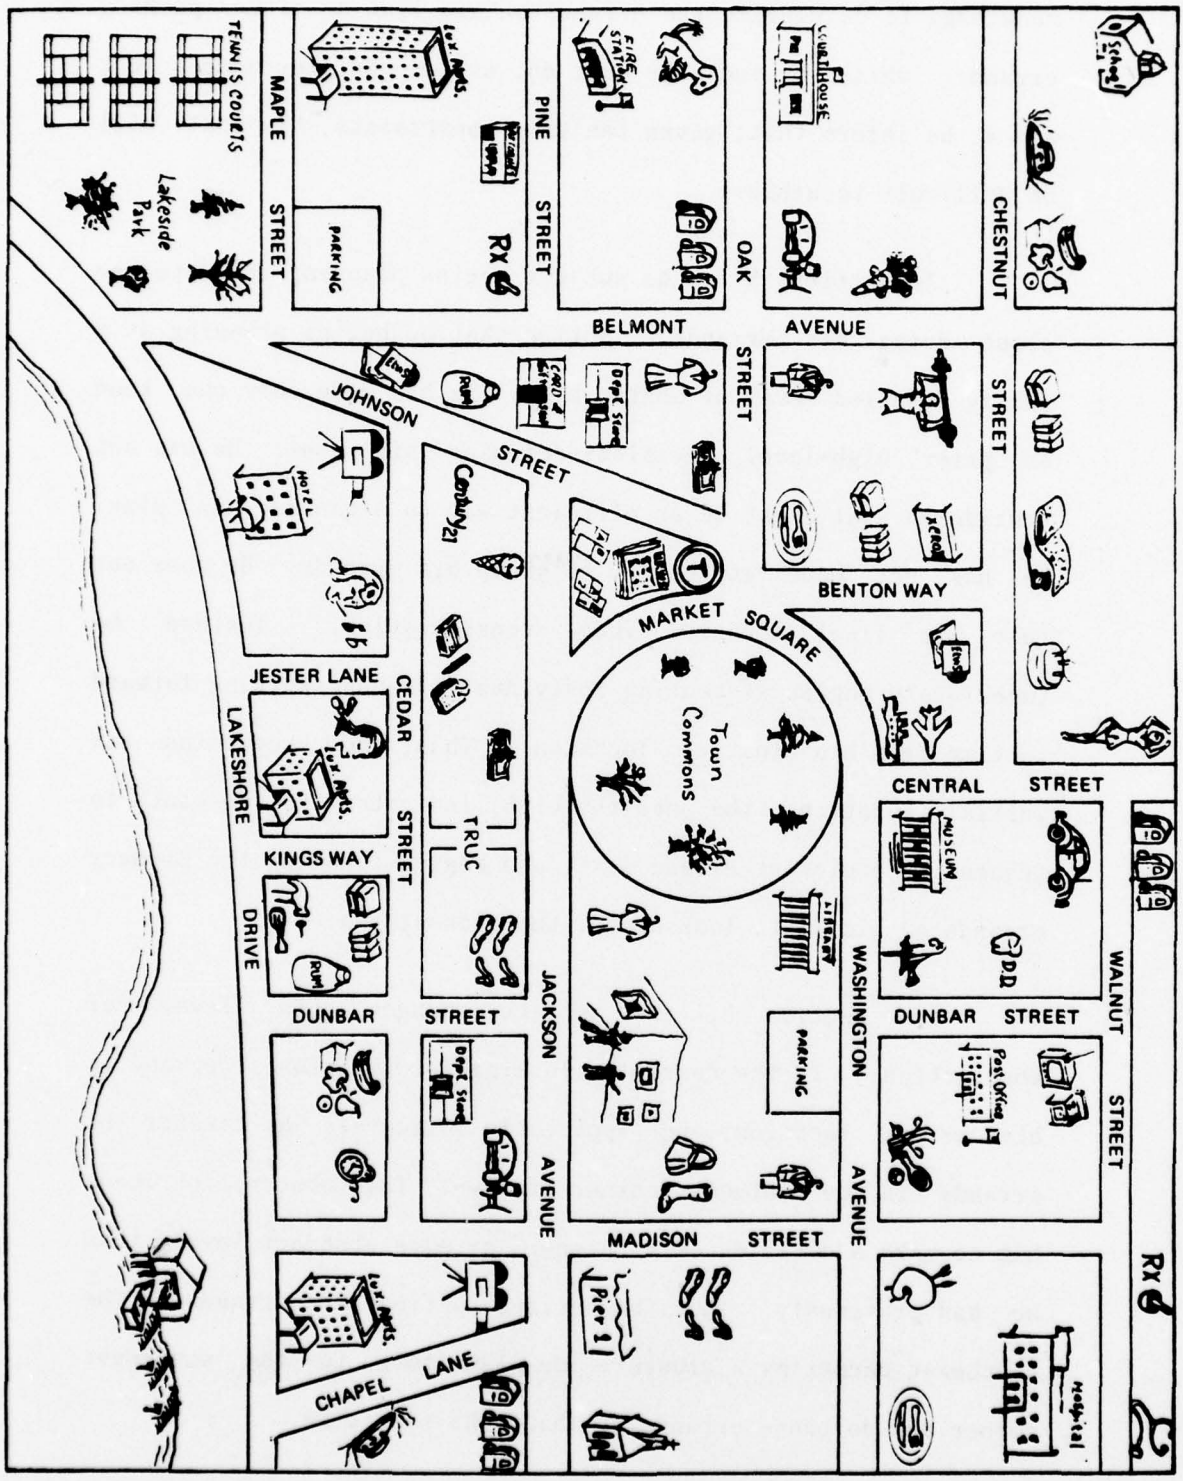

Figure 2. Town Map for the Errand-Planning Task.

In sections 1-4, the subject defines his goal and characterizes his task. Thus, in 1 and 3, he uses world knowledge to categorize the errands on his list as either primary errands, which he feels he must do, or secondary errands. In 2 and 4, he infers that, given the time constraints, his goal will be difficult to achieve.

In sections 5-7, the subject begins planning how to go about doing his errands. Notice that he begins planning at a fairly detailed level of abstraction. He has made only one kind of prior high-level decision--defining his goal. He has not considered what might be an efficient way to organize his plan. He has not made any effort to group his errands. He does not take his final location into consideration. Instead he immediately begins sequencing individual errands, working forward in time from his initial location. Thus, he ascertains his initial location, the health club, indicates that he wants to sequence the closest errand next, and begins locating the primary errands on his list, looking for the closest one.

In section 8, the subject changes his level of abstraction. In the course of looking for the closest errand to his current location, he apparently discovers a cluster of errands in the southeast corner of town. This observation leads him to make a decision at a "higher" or more abstract level than he had previously. Thus, he decides to treat the errands in the southeast corner as a cluster. He plans to go to the southeast corner and do those errands at about the same time.

In section 9, the subject modifies his high-level cluster. He discovers that one of the errands in the cluster, the movie, can also be done on the west side of town, near his final destination, the Maple Street parking structure. He changes back to the more detailed level of abstraction. Planning backward in time from his final location, he decides to end his day by going to the movie and then picking up his car. In so doing, he removes the movie from the high-level cluster.

In section 10, the subject begins to instantiate his high-level plan to go to the southeast corner at the lower, errand-sequencing level. Again, he is looking for the closest errand on his way, and he chooses the vet.

At that point, the experimenter interrupts to point out to the subject that he has overlooked several closer errands.

In sections 11 and 12, the subject incorporates the new information into his planning. His first reaction, in 11, is to continue working at the errand-sequencing level, simply considering the newly identified errands among those he might do next. However, additional observation at this level leads him to make a decision at the more abstract level. Again, he decides to treat a group of errands, those in the northwest corner of town, as a cluster. This leads him to revise his high-level plan to include two clusters of errands, the northwest cluster and the southeast cluster.

In section 13, the subject begins instantiating his new high-level plan. He notes the initial time, 11:00, and the

presence of a restaurant, another errand in the northwest cluster. These observations lead him to formulate an intermediate level plan regarding how to sequence errands within the northwest cluster. He decides to sequence the errands in that cluster to permit him to arrive at the restaurant in time for lunch.

In sections 14-15, the subject works on instantiating his revised high-level plan at a very detailed level of abstraction. Here, he not only sequences individual errands (the florist and the grocery), he specifies the exact routes he will take among them. In addition, the subject mentally simulates execution of his plan in progress, estimating how long each errand should take and computing the "current" time at each stage of the plan.

In section 16, the subject's mental simulation suggests the inference that his groceries will perish if he picks them up early in the day. This leads him to revise his low-level plan, assigning the grocery a sequential position at the end of the plan.

In section 17, the subject decides to abandon his two-cluster high-level plan in favor of his original high-level plan including only the southeast cluster. Presumably he decided that, without the grocery, there were not enough errands in the northwest cluster to occupy him until lunch.

In section 18, the subject begins instantiating his original high-level plan at a more detailed level. Again, he sequences individual errands (the florist and the vet) and

specifies exact routes among them, mentally simulating execution of his plan as he formulates it.

In sections 19-23, the subject continues working at the lowest level of abstraction. He works on planning his route from the sequenced errands to the southeast corner, mentally simulating execution of his plan in the process. In so doing, he notices a "short-cut" through the card and gift shop and incorporates it into his plan, later replacing it with one through the department store. He then notices that taking the short-cut will put him very near the newsstand. Although the newsstand is a secondary errand, he decides to incorporate it in his plan because it is so convenient. Thus, a decision at the lowest level of abstraction leads him to make a decision at the next higher level. Note also that this decision implies addition of the newsstand to the subject's definition of his goal.

In sections 24-26, the subject continues working at a low level of abstraction. He notices that his high-level plan does not include any provision for the appliance store, a primary errand. He plans to go there directly, temporarily ignoring his high-level plan to go to the southeast corner. He also notices that another secondary errand, the pet store, is on the way to the appliance store and, because it is so convenient, incorporates that errand into his plan. Again, he plans at the level of sequencing errands and specifying routes and simulates execution of the plan as he goes along. Note that these decisions imply addition of the pet store to the subject's definition of the goal. (Note also that, while the short-cut

planned in 19-23 was a short-cut to the southeast corner, it is a detour in the planned route to the appliance store.)

In sections 27-28, the subject continues working at a low level of abstraction. He notes the time, 11:50, and that lunch at the restaurant is a pending errand. He also notices a restaurant quite near the appliance store and plans his route to that restaurant.

In sections 29-32, the subject reviews the time, 12:55, and his remaining errands. He reviews his previously planned final sequence (the movie, the grocery, and the parking structure). He enumerates the primary errands remaining to be planned (the bookstore and the watch repair).

In sections 33-37, the subject continues planning at a low level of abstraction. He sequences the pending errands, using his earlier strategy of going to the closer of the two errands first. He specifies exact routes and continues to execute his plan mentally as he goes along. In planning this sequence, he notices that he must pass quite near two luxury apartments. Because visiting two luxury apartments is a secondary errand, he incorporates a visit to each apartment at the most convenient point in his plan, implicitly amending his goal to include this errand. Thus, the subject finally plans to arrive at and perform the errands in the southeast cluster. Note, however, that this occurred as a consequence of sequential planning at a low level of abstraction, rather than as a consequence of his having deliberately instantiated a high-level

plan at a low level of abstraction.

In sections 38-40, the subject evaluates his current status. He notes the time, 2:00, that he has nothing left to plan before the movie, and that the movie cannot be scheduled for an hour. He criticizes his plan for the wasted hour and considers how to minimize the wasted time. He relaxes his constraint on when to go to the grocery and decides to do that next.

In sections 41-45, the subject continues working at a low level of abstraction. He sequences his remaining errands (the grocery, the movie, the parking structure), specifying routes and simulating execution of his plan as he goes along. He notes that his plan still contains twenty-five wasted minutes and that he accomplished more than he thought he could in the time available. He resigns himself to the twenty-five minute empty period.

In sections 46-47, the subject evaluates his plan. He notes that he accomplished all of the errands on the list. He notes again that he wasted twenty-five minutes and criticizes his plan on that account.

This protocol illustrates a number of the points made above. First, the subject's plan develops incrementally at various points in the planning space we defined. He plans actions at various points in the plan's temporal sequence, and he also plans at different levels of abstraction. Second, the subject appears to plan opportunistically, "jumping about" in the planning space to develop promising aspects of the plan in

progress. For example, the planner does not plan strictly forward in time. Instead, he plans temporally-anchored sub-plans at arbitrary points on the time dimension and eventually concatenates the sub-plans. Similarly, the planner does not plan in a systematic top-down fashion across the different levels of abstraction. He frequently plans low-level sequences of errands or routes in the absence, and sometimes in violation, of a prescriptive high-level plan. Finally, decisions at a given point in the planning space appear to influence subsequent decisions at both later and earlier points in the temporal sequence and at both higher and lower levels of abstraction. The protocol exhibits examples of each of these kinds of influence.

The protocol illustrates another important component of the planning process--the ability to mentally simulate execution of a plan and to use the results of the simulation to guide subsequent planning. Mental simulation answers a variety of questions for the subject: At what time will I arrive at (or leave) a particular destination? How long will I take to perform a certain action? What sequence of operations will I perform to satisfy a particular sub-goal? How long will it take to execute a plan or partial plan? What consequences will my actions produce? What have I accomplished so far? The subject can use this information to evaluate and revise prior planning and to constrain subsequent planning.

The subject performs two kinds of mental simulation. Sometimes he simulates his plan by mentally "stepping through" a sequence of time units for each planned action (e.g., walking,

carrying a package, performing an errand). With each successive step, he extrapolates the results of each planned action, updating his understanding of the "current state" accordingly. At other times, the subject performs "event-driven" simulation. In this case, he mentally moves directly from one planned situation to another, often "ignoring" intervening actions. He then computes certain consequences arising from the transition.

More importantly, in the present context, the subject simulates execution of plans at different levels of abstraction. Thus, in sections 14-15, he simulates execution of a detailed plan. By stepping through his plan, the subject computes expected times for performing individual errands and traveling specific routes. In sections 24-26, the subject simulates execution of his high-level plan for performing errands in the northwest and then those in the southeast. Here, he performs event-driven simulation, inferring that if he executes his high-level plan, proceeding directly to the southeast corner of town, he will neglect a primary errand.

In the next section, we describe the proposed planning model in detail. The model postulates specific levels of abstraction and a structural organization for the planning space. In addition, it postulates a number of plausible planning specialists. Finally, the model embodies decision mechanisms that permit theoretical interpretation of subjects' apparently chaotic progress through the planning space.

AN OPPORTUNISTIC MODEL OF PLANNING

The proposed model assumes that the planning process comprises the independent and asynchronous operation of many distinct specialists (akin to demons in Selfridge's (1959) Pandemonium model). Each specialist makes tentative decisions for incorporation into a tentative plan. All specialists record their decisions in a common data structure, called the blackboard. The blackboard enables the specialists to interact and communicate. Each specialist can retrieve decisions of interest from the blackboard, regardless of which specialists recorded them. A specialist can combine earlier decisions with its own decisionmaking heuristics to generate new decisions. The model partitions the blackboard into several planes containing conceptually different categories of decisions. Each plane contains several levels of abstraction of the planning space. Most specialists deal with information that occurs at only a few levels of particular planes of the blackboard. Finally, specialists also establish linkages on the blackboard to reflect causal or logical relationships among various decisions.

The proposed model generalizes the theoretical architecture developed by Reddy and his associates (Cf. CMU Computer Science Research Group, 1977; Lesser, Fennell, Erman, & Reddy, 1975; Erman & Lesser, 1975; Lesser & Erman, 1977; Hayes-Roth & Lesser, 1977) to enable computers to perform complex problem-solving tasks. This architecture was conceived for the Hearsay-II speech-understanding system. Others have since applied it to image understanding (Prager, Nagin, Kohler, Hanson,

& Riseman, 1977), reading comprehension (Rumelhart, 1976), protein-crystallographic analysis (Nii & Feigenbaum, 1977), and inductive inference (Soloway & Riseman, 1977). The architecture's rapid acceptance reflects its versatility as a model of the problem-solving processes involved in "interpretation" tasks. The model described below is, to our knowledge, the first attempt to adapt the Hearsay-II architecture to a "generation" problem.

We organize our discussion of the details of the planning model around the structure and content of the blackboard. The blackboard comprises five planes (see Figure 3), each of which represents conceptually different categories of decisions (see also Engelmore & Nii, 1977). We have already characterized the plan plane in our discussion of the protocol. Plan decisions indicate actions the planner intends to take in the world. We characterize the other four planes briefly as follows. Decisions on the plan-abstractions plane characterize desired attributes of potential plan decisions. The knowledge-base plane contains observations and computations about relationships in the world that bear on the planning process. The meta-plan plane contains higher-level decisions regarding how the planner intends to approach the planning problem itself. Finally, the executive plane contains decisions about how the planner intends to allocate his cognitive resources among the other four planes during the planning process. In the remainder of this section, we discuss the individual planes and their constituent levels of abstraction in more detail. We also explicate the behavior of

several illustrative specialists.

Meta-plan decisions indicate what the planner intends to do during the planning process. This plane has four levels. Beginning at the top, the problem definition describes the planner's conception of the task. It includes descriptions of the goal, available resources, possible actions, and constraints. In the errand-planning task, for example, the problem definition would include the list of errands, contextual information, and associated instructions. The problem-solving model indicates how the planner intends to represent the problem and generate potential solutions. For example, the planner might view the errand-planning task as an instance of the familiar "traveling salesman" problem (Christophides, 1975) and approach the problem accordingly. Problem-solving models can also consist of general problem-solving strategies, such as "divide and conquer," "define and successively refine" (Cf., Aho, Hopcroft, & Ullman, 1974), etc. Policies specify general criteria the planner wishes to impose on his problem solution. For example, the planner might decide that his plan must be efficient or that it should minimize certain risks. Solution-evaluation criteria indicate how the planner intends to evaluate prospective plans. For example, he might decide to speculate on what could go wrong during execution and insure that his plan is robust over those contingencies.

Plan decisions indicate actions the planner actually intends to take in the world. Decisions at the four levels form a potential hierarchy, with decisions at each level specifying a more refined plan than those at the next higher level. Beginning

at the most abstract level, outcomes indicate what the planner intends to accomplish by executing the finished plan. In the errand-planning task, for example, outcomes indicate what errands the planner intends to accomplish by executing the plan. Designs characterize the general approach by which the planner intends to achieve the outcomes. For the errand-planning task, designs characterize the general route the planner intends to take to accomplish the intended errands. Procedures specify specific sequences of molar actions. Thus, for the errand-planning task, procedures specify sequences of errands. Operations specify sequences of molecular actions. In the errand-planning task, operations specify the route by which the planner will proceed from one errand to the next.

Plan-abstraction decisions characterize desired attributes of potential plans. These abstract decisions serve as heuristic aids to the planning process suggesting potentially useful qualities of planned actions. Each level of the plan-abstraction plane characterizes types of decisions suggested for incorporation into the corresponding level of the plan plane. For example, the planner might indicate an intention to do all of the "critical" errands. This intention could stimulate efforts to partition the errands into critical and non-critical sets. At a lower level, he might generate a scheme to fabricate a design employing gross spatial clusters of errands. This scheme might motivate a search for coherent clusters. At the next level, he might develop a strategy suggesting that errands in the current cluster be completed before moving on to errands in another

cluster. This strategy would presumably constrain procedural sequences eventually incorporated into the plan. Finally, he might adopt a tactic that suggested searching for a short-cut between one errand and the next. This tactic might lead to the discovery and use of one particular short-cut.

The knowledge base records observations and computations about relationships in the world which the planner generates while planning. This knowledge supports two types of planning functions: situation assessment, the analysis of the "current state" of affairs; and plan evaluation, the analysis of the likely consequences of hypothesized actions. Again, the levels of the knowledge base form a hierarchy and correspond to the levels of the plan and plan-abstraction planes. Each level of the knowledge base contains observations and computations useful in instantiating decisions at the corresponding level of the plan-abstraction plane or generating decisions at the corresponding level of the plan plane. Because the levels of the knowledge base contain problem-specific information, we have given them problem-specific names. At the errand level, for example, the planner might compute the time required to perform all of the currently intended errands to evaluate the plan's gross feasibility. At the layout level, he might observe that several errands form a convenient spatial cluster and, as a consequence, formulate a design organized around clusters. At the neighbor level, the planner might observe that two planned errands are near one another and, as a consequence, adopt a procedural decision sequencing those two errands. At the route

level, he might detect a previously unnoticed short-cut and then exploit it in an operation-level route between two planned errands.

In addition to the abstractness dimension, the plan, plan-abstraction, and knowledge-base planes have a second dimension corresponding to the time period spanned by proposed decisions. In addition, suitable blackboard representations exist for recording decisions about simultaneous and event-contingent actions and for recording competing alternative decisions.

Before describing the executive plane of the planning blackboard, we must discuss planning specialists. Specialists generate tentative decisions for incorporation into the plan in progress. Decisions become final only after the planner has accepted an overall plan. This ordinarily requires that he has formulated a complete plan and determined that it satisfies solution evaluation criteria recorded on the meta-plan plane.

Most specialists work with decisions at only two levels of the blackboard. One level contains decisions (previously generated by other specialists) that stimulate the specialist's behavior. The other is the level at which the specialist records its own modifications to the blackboard. The circle and arrow ends of the arc associated with each specialist in Figure 3 indicate these two levels respectively. For example, the "strategist" (on the plan-abstraction plane) responds to prior scheme decisions by generating strategies useful in implementing

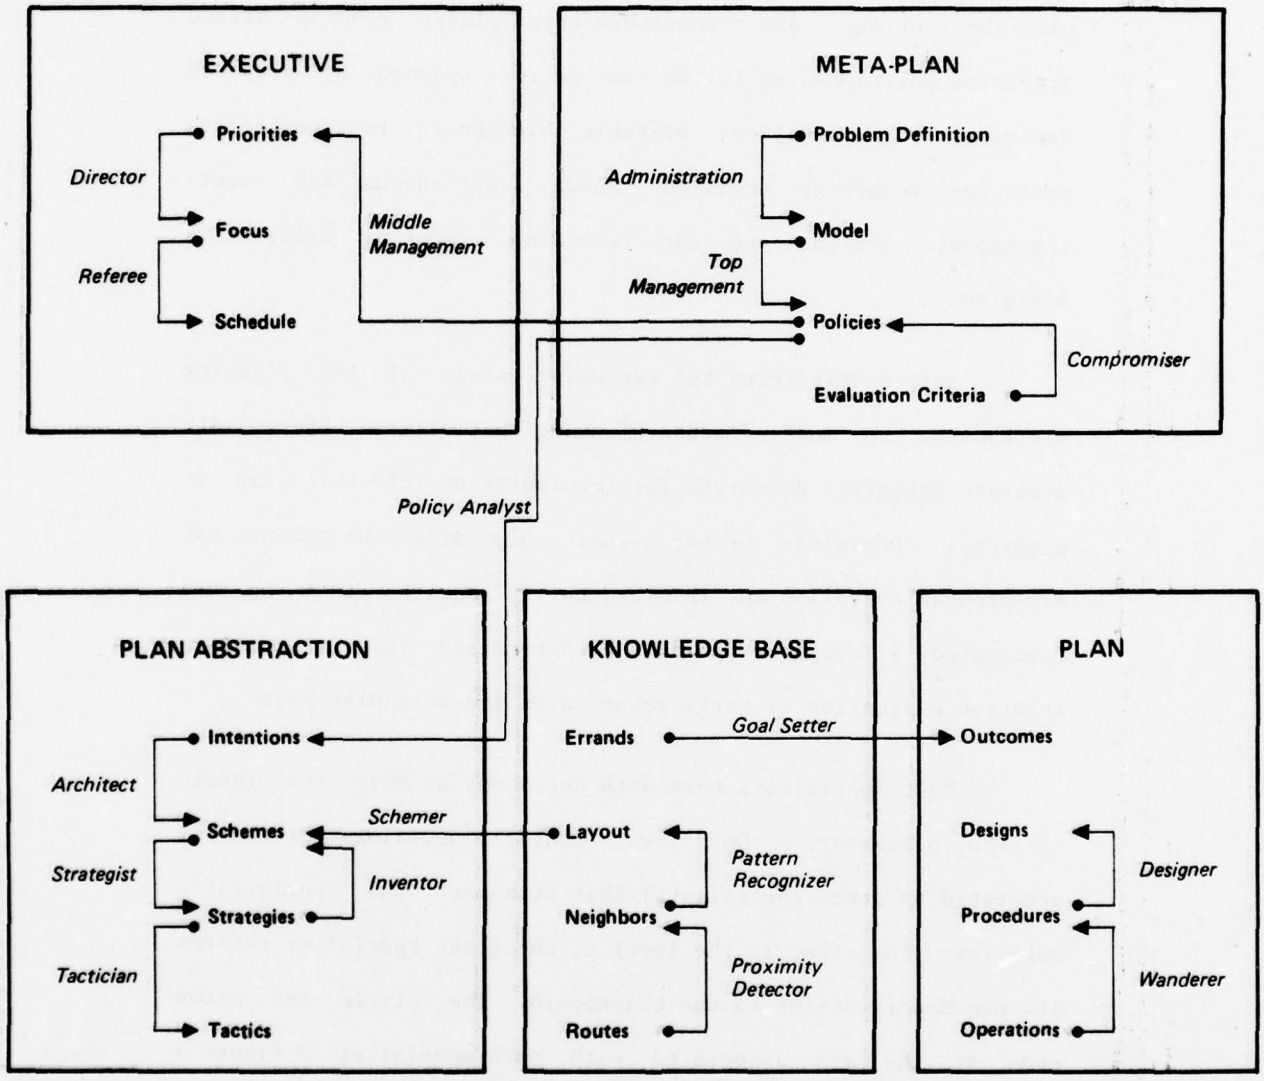

Figure 3. The Planning Blackboard and the Actions of Illustrative Specialists.

those schemes. Suppose, for example, one specialist had generated a scheme to travel around among spatial clusters of errands, doing the errands in one cluster before moving on to the next. The strategist would generate a strategy for sequencing individual errands according to this scheme. One such strategy would be to perform all pending errands in the current cluster before performing errands in any other cluster.

Note that the arcs in Figure 3 indicate that both bottom-up and top-down processing occur and that the two levels indicated by an arc need not be adjacent or even on the same plane of the planning blackboard.

The theory operationalizes specialists as condition-action modules. The condition component of a specialist characterizes decisions whose occurrences on the blackboard warrant a response by the specialist. The occurrence of any of these decisions invokes the specialist. For example, the occurrence of a new scheme on the plan-abstraction plane invokes the strategist. The action of a specialist module defines its behavior. For example, the strategist generates strategies for implementing designs. In addition to recording new decisions, each specialist records relational linkages among the decisions with which it deals. For example, the strategist records support linkages connecting the scheme decision that invokes it to the strategies it generates for implementing that design.

We have selected the specialists shown in Figure 3 for

illustrative purposes. We have excluded many other possible specialists for simplicity. The mnemonic names of the specialists and the preceding discussion of levels make most of the specialists self-explanatory. Therefore, we describe only a few of them here.

Operating within the plan-abstraction plane, the "architect" responds to intentions by generating a scheme for a design. In the errands task, for example, the architect might respond to an intention to do all the important errands by generating a scheme to travel around among spatial clusters of important errands, doing the errands in one cluster before moving on to the next.

Operating between the knowledge-base and plan-abstraction planes, the "schemer" responds to the layout of errands by suggesting an appropriate scheme. For example, the presence of one or more spatial clusters would invoke the pattern recognizer. It would respond by generating a cluster scheme.

Operating within the plan plane, the "designer" responds to a useful procedure by generating a design to exploit that procedure. For example, the designer might notice a procedure capable of accomplishing several errands in the same neighborhood in sequence. It might respond by generating a "cluster" design of the sort described above to exploit that kind of procedure.

Operating between the meta-plan and plan-abstraction planes, the "policy analyst" responds to policies by generating intentions. For example, it might respond to a policy

emphasizing efficiency by generating an intention to neglect "out-of-the-way" errands.

During planning, each of the independent specialists monitors the blackboard for the occurrences of decisions specified in its condition. Invoked specialists queue up for execution, and an executive decides which will execute its action.

We have formalized the executive as the fifth plane of the blackboard. Decisions made at the three levels on this plane form a hierarchy, with decisions at each level potentially refining ones at the level above. Starting at the top, priority decisions indicate preferences for allocating processing activity to certain areas of the planning blackboard before others. For example, given a "traveling salesman" model, the planner might decide to determine what errand sequences he could do conveniently, rather than deciding what errands he ought to do. Focus decisions indicate what kind of decision to make at a specific point in time, given the current priorities. For example, the planner might decide to focus his attention on generating an operation-level refinement of a previously generated procedure. Finally, schedule decisions indicate which of the currently invoked specialists, satisfying most of the higher-level executive decisions, to execute. If, for example, given current priorities and focus decisions, both the architect and the pattern recognizer had been invoked, the planner might decide to schedule the pattern recognizer. Schedule decisions select specialists on the basis of relative efficiency,

reliability, etc. (Hayes-Roth & Lesser, 1977).

Like the other planes of the planning blackboard, the executive plane includes decisions motivated by prior decisions on the same or other blackboards. For example, "middle management" responds to policies on the meta-plan plane by generating priorities on the executive plane. The "referee" uses focus decisions in deciding which of the currently invoked specialists to schedule. The executive plane differs from the other four planes of the planning blackboard because decisions recorded there do not motivate decisions recorded on other blackboards. Instead, they determine which invoked specialists can execute their actions on their designated planes of the blackboard.

Under the control of the executive, the planning process proceeds through successive invocation and execution of the various operational specialists. The process continues until both: (a) the planner has integrated mutually consistent decisions into a complete plan; and (b) the planner has decided that the existing plan satisfies the evaluation criteria recorded on the meta-plan plane of the blackboard.

ANALYSIS OF THE PLANNING PROTOCOL UNDER THE OPPORTUNISTIC MODEL

The opportunistic model captures the gross characteristics of the observations and decisions recorded in the thinking aloud protocol discussed above. In addition, the model accounts for each individual statement in the protocol. In this

section, we illustrate the descriptive power of the model for sections 1-10 of the protocol.

Figures 4-8 show blackboard representations of the subject's verbalizations as individual decisions. They also show how individual specialists respond to the presence of particular decisions on the blackboard by generating other decisions and recording them at appropriate locations on the blackboard. Each arrow represents the invocation and execution of a specialist. Thus, an arrow from one decision to another indicates that the former decision invoked a specialist that recorded the latter decision. In order to clarify the flow of activity, we have numbered decisions in Figures 4-8 according to their presumed order of occurrence.

We have omitted only one kind of decision from these illustrations--scheduling decisions. As discussed above, at each point in the sequence of recorded decisions, a scheduling decision selects one of the currently invoked specialists to execute its action. We have omitted these decisions from Figures 4-8 for simplicity. However, it is appropriate to assume that a scheduling decision selected each of the indicated specialist actions (noted by arrows).

Figure 4 shows the blackboard representation of sections 1-4 of the protocol. In sections 1 and 3, the subject works through the list of errands, assigning binary importance values (primary versus secondary) to each one. In sections 2 and 4, the subject remarks that the large number of primary errands implies

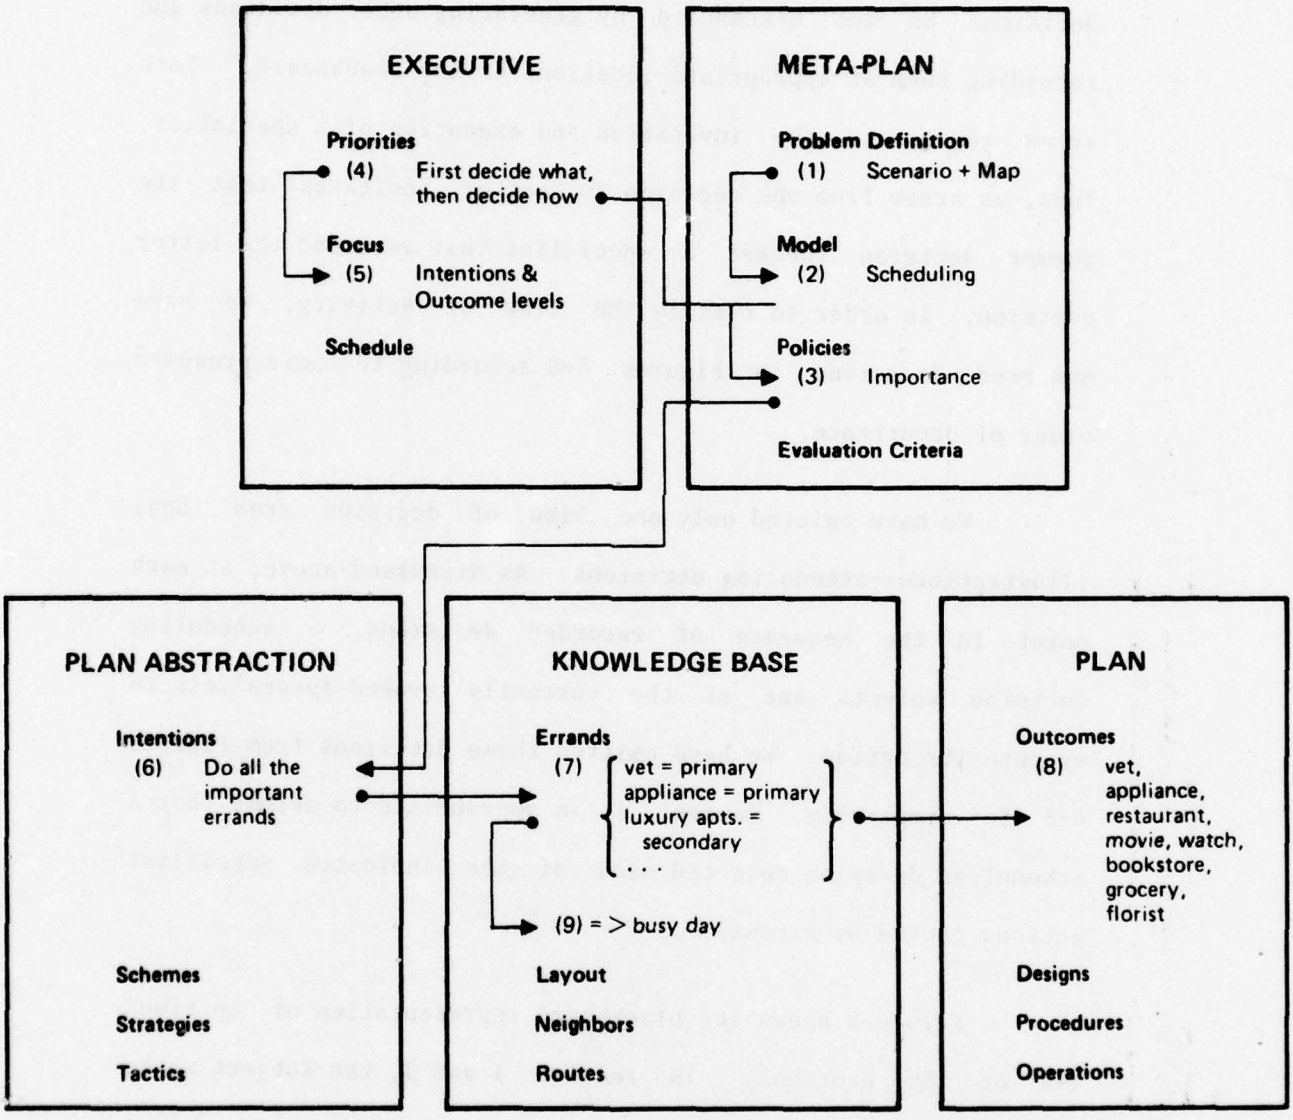

Figure 4. Blackboard Representation of Sections 1-4 of the Protocol.

that he will have a busy day. According to our assumptions, a specialist calculates importance values for individual errands and records these at the errands level of the knowledge base. However, we assume that a considerable amount of activity, unstated in the protocol, preceded and motivated this action. Figure 4 shows the blackboard representation of this implicit activity.

The subject begins the task with a problem definition (1), including the scenario and map provided by the experimenter. The protocol suggests that the subject identifies the problem as a "scheduling" problem (2). In other words, the subject apparently views the task as one in which he cannot do all of the things he wants to do and, therefore, must decide which things to do and then how to do them. The appearance of this problem-solving model on the blackboard presumably invokes two other specialists. One generates and records a useful policy (3), emphasizing the importance of individual errands. The other generates and records an appropriate set of priorities (4). The priorities, in turn, motivate a decision to focus on the intentions and outcomes levels of the plan-abstraction and plan planes (5). Given this focus and the errand-importance policy, a specialist records an intention to do all the important errands (6). This intention presumably invokes the specialist described above that calculates the errand-importance values actually stated in the protocol (7). This activity implies another unstated decision, that the intended outcomes include the designated primary errands (8). Finally, the statements in

sections 2 and 4 of the protocol imply that the errand-importance calculations invoke another specialist that infers: "It's going to be a busy day" (9).

Figure 5 shows the blackboard representation of section 5 of the protocol. In section 5, the subject states: "All right. We are now at the health club." This statement conveys a procedural specification of the initial location (13). Figure 5 shows the implicit sequence of activity that produced this statement, given the prior state of the blackboard shown in Figure 4. First, having decided what to do (8), the subject proceeds to his second priority, deciding how to do it. Accordingly, he changes focus to the lower levels of the blackboard (10). Given this focus, a strategy-generating specialist records its decision to plan forward from the initial location (11). This decision motivates another specialist to identify the initial location (12) which, in turn, motivates a specialist to record the initial location at the procedure level of the blackboard (13).

Figure 6 shows the blackboard representation of sections 6-8 of the protocol. In section 6, the subject asks, "What is going to be the closest one?" This question indicates a strategic decision to plan to perform the closest errand next in the procedural sequence (14). The appearance of this strategy on the blackboard invokes a specialist that evaluates the relative proximities of other primary errands to the initial location, the health club (15). Section 7 of the protocol describes these evaluations.

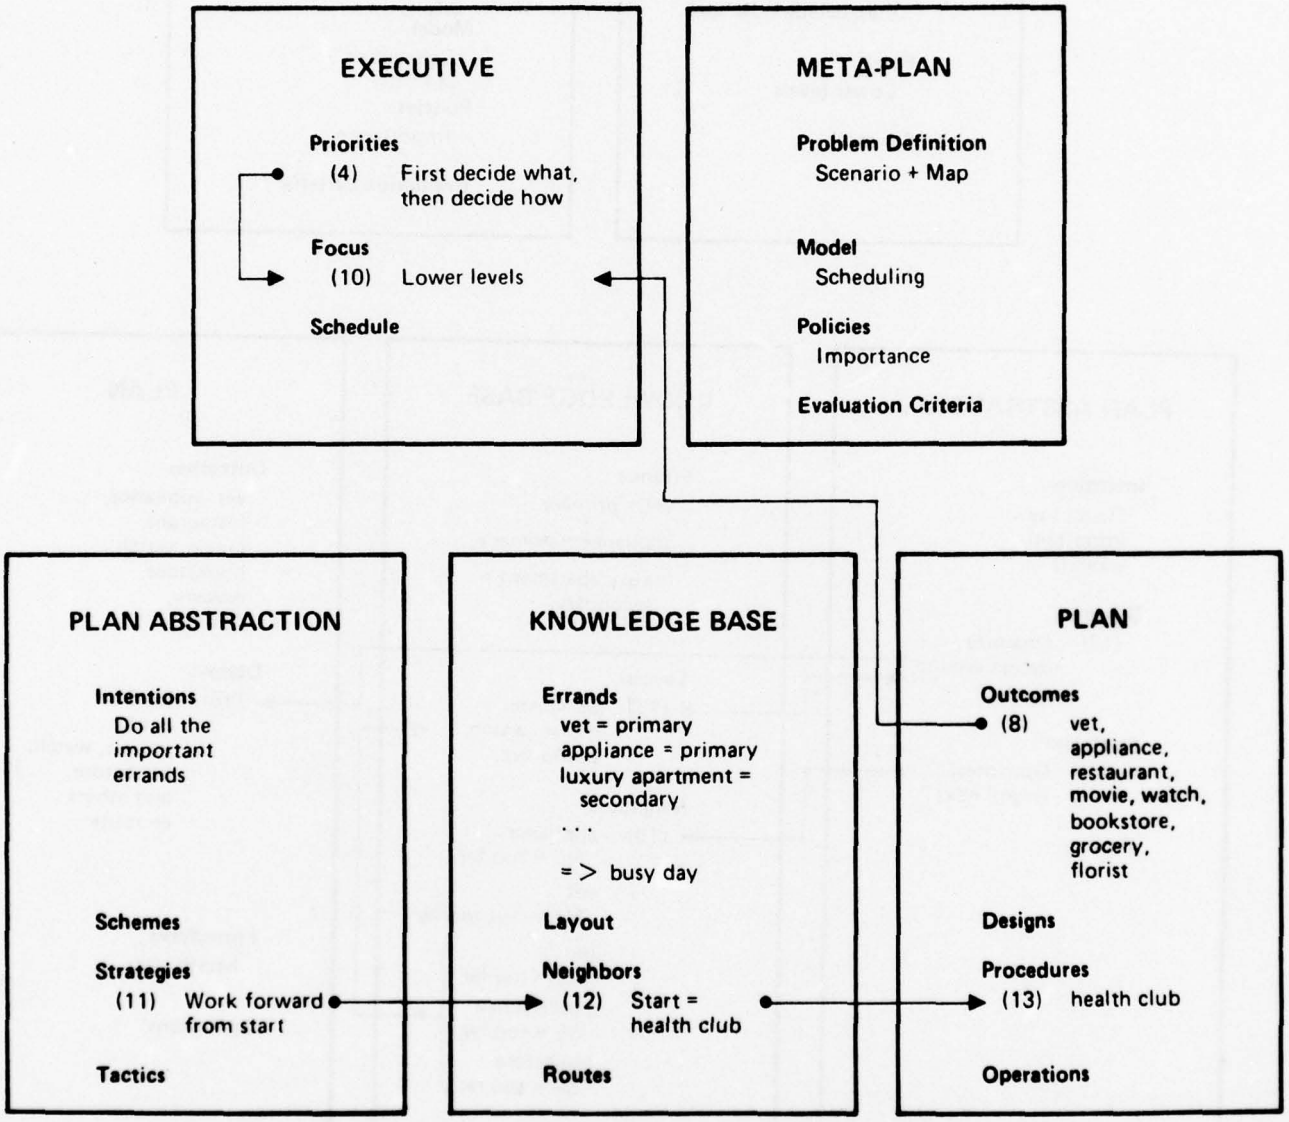

Figure 5. Blackboard Representation of Section 5 of the Protocol.

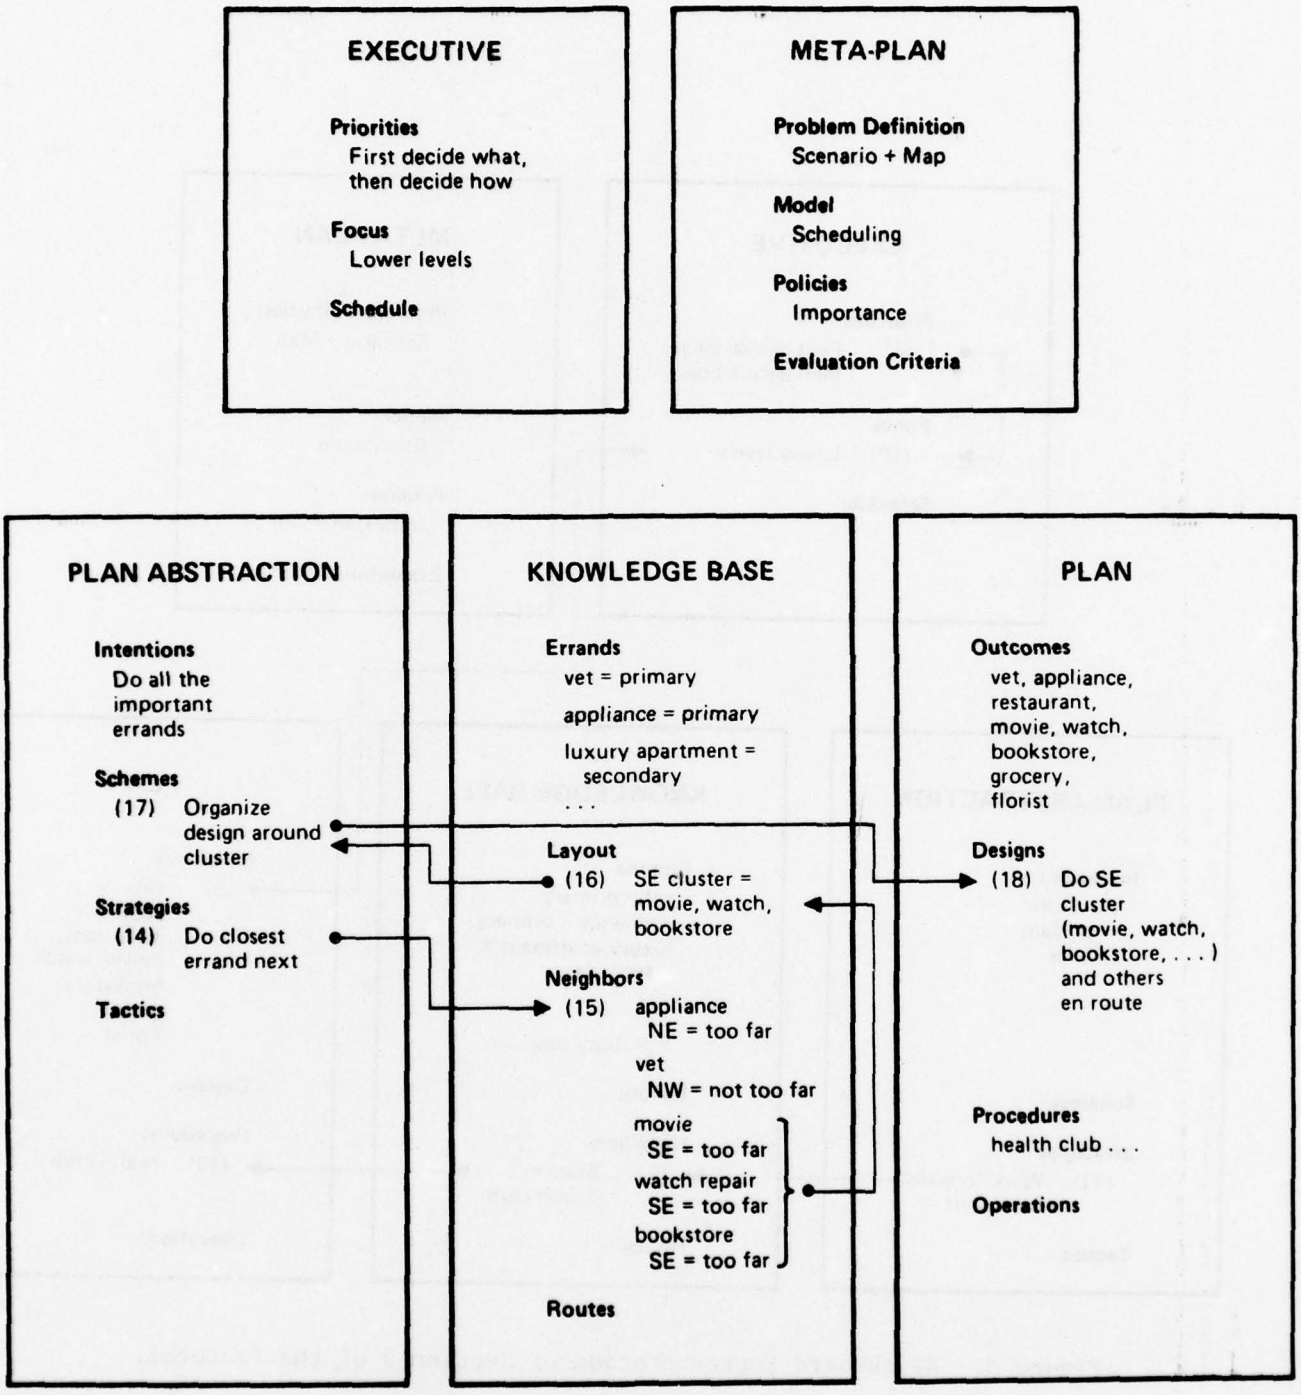

Figure 6. Blackboard Representation of Sections 6-8 of the Protocol.

Section 8 of the protocol reflects a discontinuity in the planning process. The preceding statements aim toward recording the second errand in the procedural sequence. Instead, however, the subject states in section 8: "Probably it would be best if we headed in a southeasterly direction. Start heading this way. I can see later on there are a million things I want to do in that part of town." This statement expresses a higher-level design, recorded on the blackboard as a decision to perform the errands in the southeast cluster, performing whatever other errands occur along the route from the initial location to the southeast cluster (18).

Let us consider how the subject might have arrived at this design. The subject's immediately-preceding overt activity, evaluation of proximities, requires him to locate each errand in the list. In so doing, the subject locates (at least) three consecutive errands, the movie, the watch repair, and the bookstore, in the southeast corner of town. Apparently, this sequence of observations invokes a specialist that identifies clusters of errands and records the identity of the detected cluster at the layout level of the knowledge base (16). The appearance of the cluster on the blackboard invokes another specialist that generates schemes. It suggests exploiting the spatial cluster of errands by organizing a design around it (17). Another specialist responds to the new scheme and the identified cluster by recording the appropriate design on the blackboard (18).

Figure 7 shows the blackboard representation of section 9

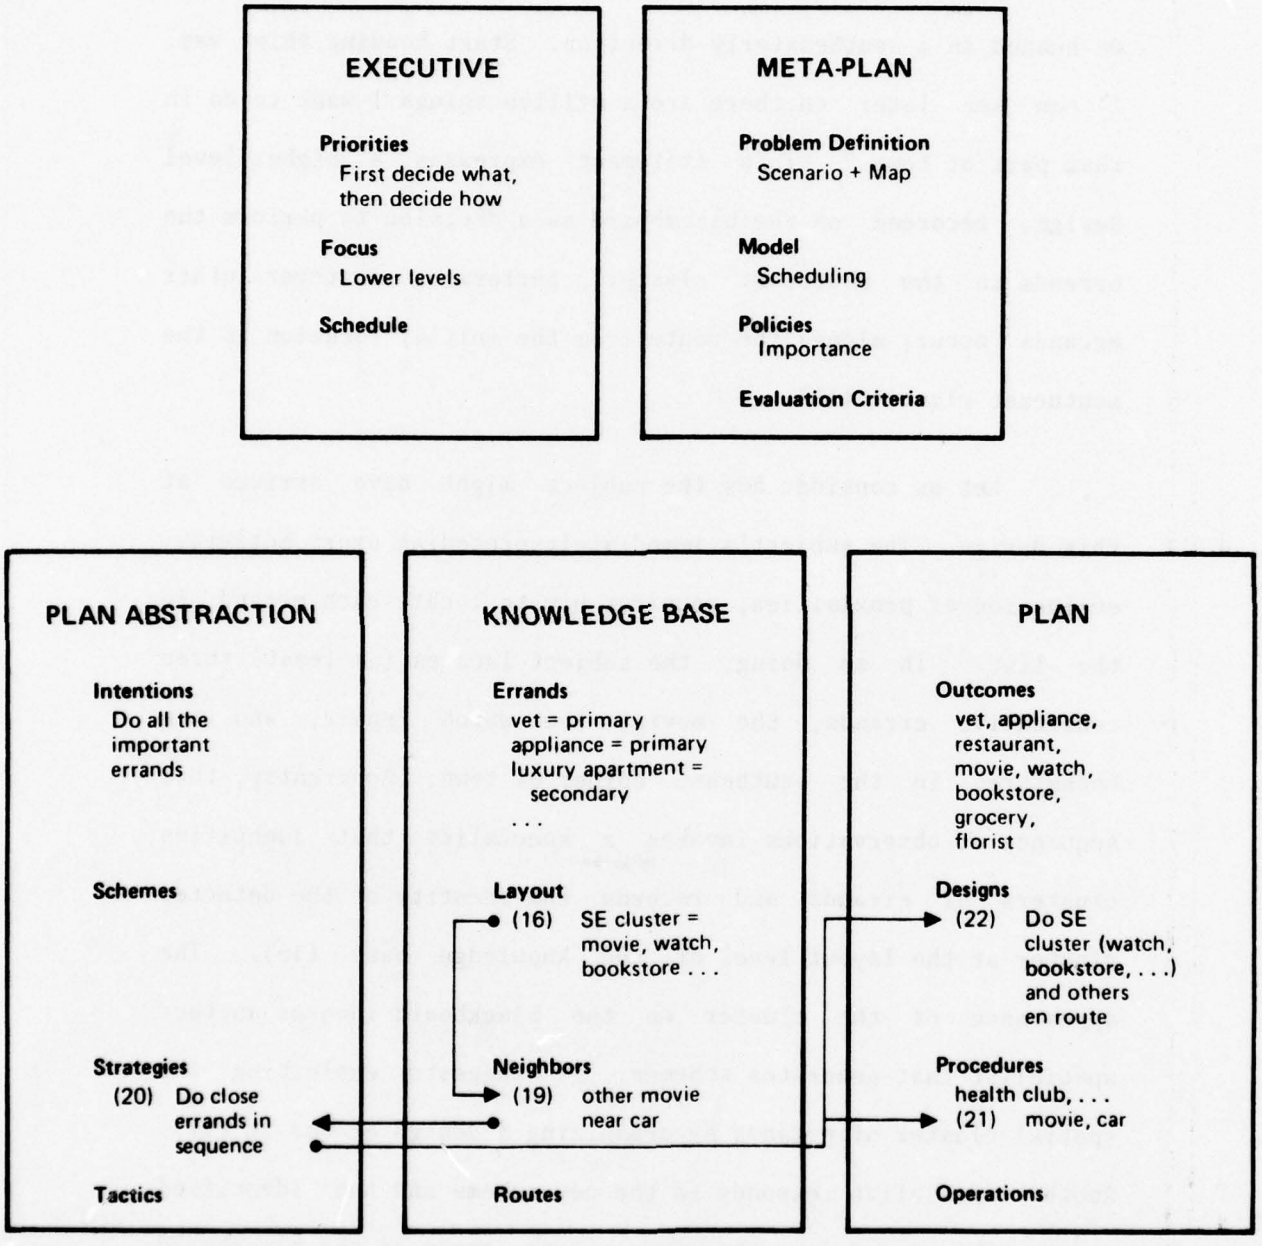

Figure 7. Blackboard Representation of Section 9 of the Protocol.

of the protocol. In section 9, the subject indicates a procedure decision to sequence the movie right before picking up his car at the end of the day (21). He tells us explicitly that, in so doing, he is removing the movie from the previously-defined southeast cluster (22). He also tells us why he has made this decision: because it would be "convenient to get out of the movie and go to the car" (20).

Figure 7 models these decisions, beginning with the subject's prior definition of the southeast cluster (16). Presumably, attention to one of the errands in the cluster, the movie, invokes a specialist that notices another movie on the west side of town close to the parking structure (19). The proximity of these two errands invokes a specialist that suggests a more general strategy to perform two proximate errands in sequence (20). This new strategy invokes another specialist that records the suggested sequence, movie-car, at the procedure level of the plan plane (21) and amends the prior design accordingly (22).

Figure 8 shows the blackboard representation of section 10 of the protocol. In section 10, the subject decides to go to the vet after the health club because it is the closest primary errand. Thus, section 10 conveys a procedure-level decision (26) and the strategy that motivated it (24). We assume that the presence of a modified design on the blackboard motivates a narrowing of the focus to aim at instantiating the design at the procedure level (23). In accordance with this focus, the design also invokes a specialist that generates a strategy to do the

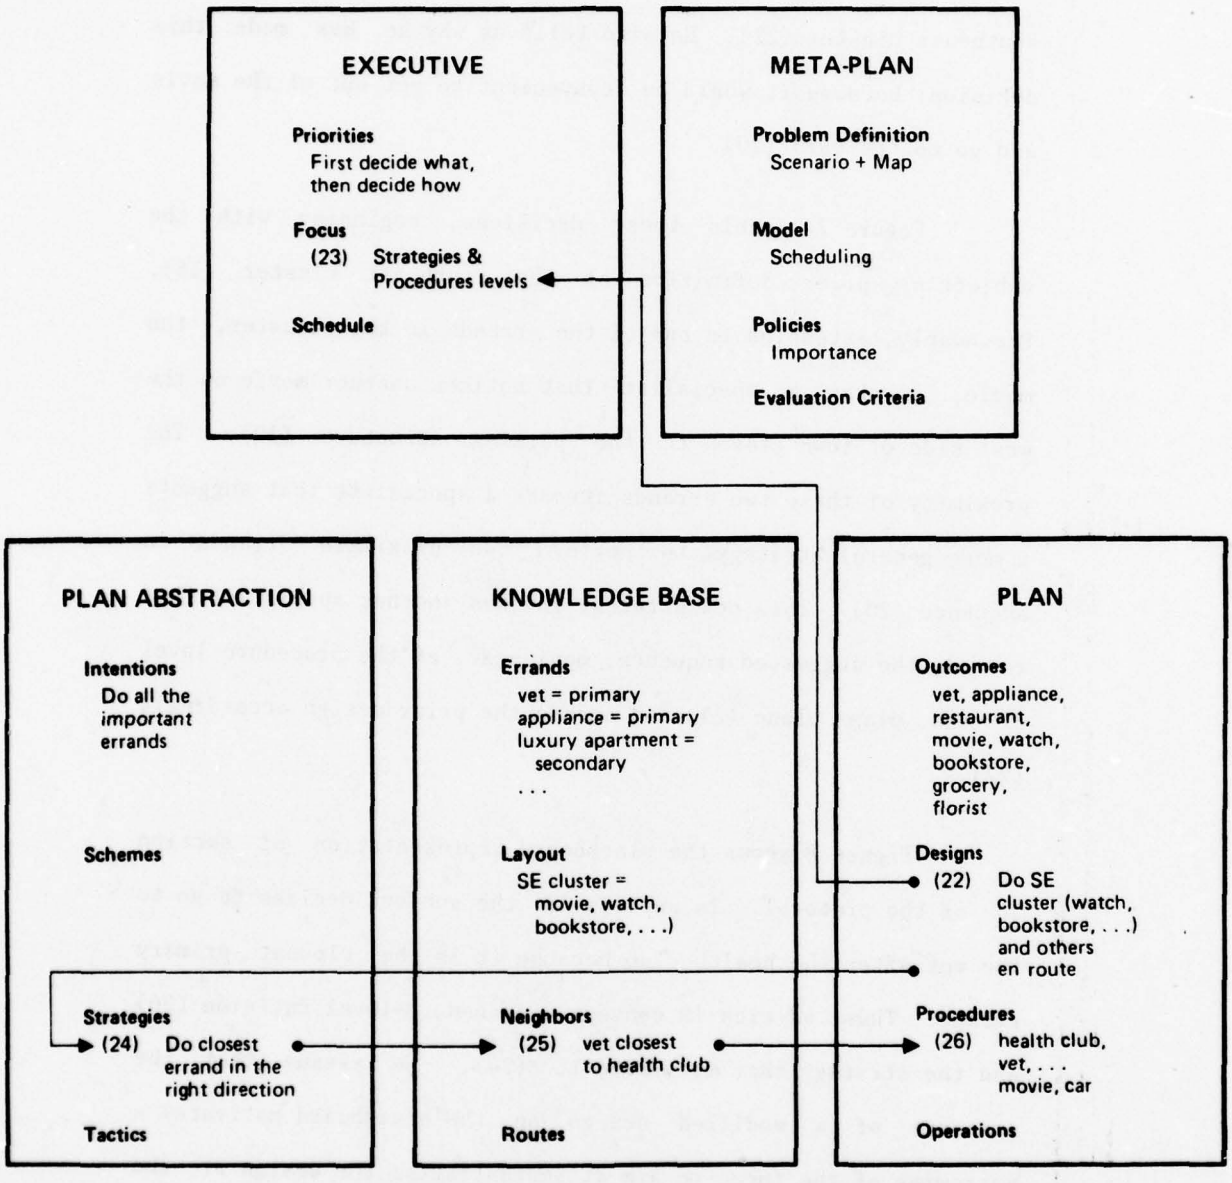

Figure 8. Blackboard Representation of Section 10 of the Protocol.

closest errand in the right direction (24). This strategy invokes a specialist that evaluates the proximities of individual errands at the neighbors level of the knowledge base (25). Finally, the observation that the vet is the closest errand to the initial location, the health club, invokes a specialist that records the vet as the next errand in the procedural sequence (26).

We can analyze the remainder of the protocol in much the same fashion. However, we conclude the analysis at this stage for brevity.

The analysis clarifies the points suggested by the informal analysis discussed earlier. The subject plans at different points in the planning space along both temporal and abstractness dimensions. In particular, the subject appears to make decisions at each of the postulated levels on all five planes of the blackboard. Further, the subject makes decisions opportunistically. Rather than working systematically through the levels along either of the two dimensions, he enters the planning space at various points and moves about freely within it. The subject's observations and computations on the available data (the map and the scenario) exert a powerful influence on the point in the planning space at which he makes each successive decision. This indicates a strong "bottom-up" component to the planning process. However, prior decisions at both higher and lower levels influence the subject's decisions, as predicted by the model.

COMPARISON WITH EARLIER MODELS OF PLANNING

Miller, Galanter, and Pribram (1960) provide a convenient starting point for our discussion of previous models of planning. Their work emphasized the importance of plans as the "guiding force" behind all human behavior. In addition, they argued persuasively for efforts to understand behavior at each of the many possible levels of analysis (abstraction). Accordingly, they defined a plan as "any hierarchical process in the organism that can control the order in which a sequence of operations is to be performed" (p. 16). These authors had little to say about how plans get formed. Instead, they focused on the execution of existing plans to control behavior. Exploiting the presumed hierarchical structure of plans, they suggested a strictly "top-down" execution procedure. That is, they assumed that people execute plans and generate behavior by successively refining abstract plans (high in the hierarchy) into more detailed plans (low in the hierarchy). The lowest-level plan constitutes a sequence of mental or physical operations.

The conception of plans as abstract representations of problem solutions appears in much of the problem-solving literature. For example, Newell, Shaw, and Simon (1963) discussed planning as a three-step sequence consisting of (1) simplification of a problem to omit details, (2) solution of this more general problem, (3) and refinement of the solution back into the detailed context. Several researchers (Greeno, 1974; Newell, & Simon, 1972; Reed, Ernst, & Banerji, 1974; Thomas, 1974) have incorporated this view of planning into their accounts

of human problem-solving. They assume that problem-solvers progress through a series of "cognitive states," arriving at each one through application of a sequence of moves prescribed by a particular strategy. Atwood and Polson (1976) and Jeffries, Polson, Razran and Atwood (1977) adopt a similar definition of planning.

More recently, Sacerdoti (1975) has implemented a computer program that plans by successive refinement. His program, NOAH, formulates problems in terms of high-level goals that specify sequences of actions (for example, the monkey should get the bananas and then eat them). NOAH expands each constituent sub-goal into additional subgoals, maintaining any indeterminate sequential orderings as long as possible. In this manner, NOAH eventually generates correct plans specifying sequences of elementary actions. When executed, these actions transform initial conditions into a series of intermediate conditions, culminating in the goal state. (See also: Ernst & Newell, 1969; Fahlman, 1974; Fikes, 1977; Fikes & Nilsson, 1971; Sacerdoti, 1974; Sussman, 1973.)

This view of planning as simplification, problem-solving, and successive refinement differs from the proposed model in three important ways. In the subsequent paragraphs, we discuss these differences and relate them to the protocol discussed above.

Top-Down versus Multi-Directional Processing

While the earlier work assumes that planning is a top-

down process, the proposed model characterizes planning as a multi-directional process. The diverse observations people make while planning often guide subsequent planning. Some of these observations arise from planning at an abstract level and guide subsequent planning at a more detailed level. The errand-planning protocol illustrates this kind of top-down processing in section 10, where the subject begins to instantiate a previously planned design at the lower procedure level. However, observations also arise from planning at a low level and guide subsequent planning at a more abstract level. The protocol illustrates this kind of bottom-up processing in section 8 where the subject formulates a design based on observations related to previous decisions at the lower procedure level. Many other examples of both top-down and bottom-up processing appear throughout this protocol and the others we have collected.

The sample protocol supports the multi-directional assumption over the top-down assumption in another way. If the subject were operating in a top-down fashion, he would begin planning at the highest (most abstract) level of the planning space. He could plan at a lower level only if he had already planned that particular sub-task at all higher levels. The errand-planning protocol disconfirms this prediction repeatedly. The subject begins forming his actual plan at a relatively low level, the procedure level. Thus, he plans at this level in the absence of any corresponding high-level plans. Similar instances of planning a sub-task at a low level without having previously planned it at higher levels appear throughout this protocol and

our others. These findings follow directly from the multi-directional assumption.

Complete versus Opportunistic Planning

A second difference between the earlier view of planning and the proposed model concerns the relative completeness attributed to abstract plans. The earlier work assumes that, while initial plans may be abstract, they will be complete and fully integrated. Under a breadth-first processing assumption, this requires that complete plans at each level must precede any planning at the next lower level. Under a depth-first processing assumption, it requires only that the highest-level plan must be complete before planning activity can proceed at lower levels. Under both assumptions, the earlier view presupposes that complete plans will eventually exist at all levels of abstraction.

By contrast, we assume that planning is opportunistic and, therefore, will rarely produce complete plans in the systematic fashion described above. We assume that people make tentative decisions without the requirement that each one fit into a current, completely integrated plan. As the planner relates each new decision to some subset of his previous decisions, the plan grows incrementally. Further, the developing plan need not grow as a single integrated plan. Various sub-plans can develop independently either within or between levels of abstraction. The planner can incorporate these sub-plans into his final plan as he wishes.

The sample protocol provides evidence for these assumptions. For example, in section 9, having established only his initial location at the procedure level, the subject plans a sequence of two errands with which to conclude. In the following several sections of the protocol, he intermittently plans alternative designs (none of which covers the planned concluding sequence) and initial sequences of errands (none of which he concatenates with the planned concluding sequence). Similar partial plans appear throughout the protocol as well as in the other protocols we collected. These findings confirm our assumption that specialists record tentative decisions in various locations on the blackboard in response to relevant prior decisions.

Hierarchical versus Heterarchical Plan Structures

Earlier conceptions of plans as hierarchical structures responded to the appealing simplicity of hierarchically structured programs and the successive refinement method. None of our observations denies the putative merits of these hierarchical approaches. Of course, one can always interpret a sequence of actions as a hierarchy with some number of levels. Therefore, one must perform some more informative analysis to contrast hypothesized hierarchical plans with more complex plan structures. More importantly, a satisfactory theory of planning must describe all decisions made during the planning process as well as those that appear in completed plans.

Our efforts to model the planning process suggest that

people make many decisions that do not fit a simple hierarchical structure. Under the proposed model, one might attempt to construe the final set of decisions on the plan plane as a hierarchical structure, but our protocols do not provide strong evidence for a such a structure. For example, the design maintained throughout most of the sample protocol dictates that errands on the way to the southeast cluster should be performed first followed by those errands within the cluster itself. However, much of the subject's planning at lower levels concerns errands not covered by this design (e.g., the newsstand, the pet store, the appliance store, and the restaurant).

The assumption of hierarchical plan structure becomes more tenuous if we consider the many other kinds of decisions our subject made while planning. We have observed four categories of decisions that do not describe what the subject actually plans to do at all. These correspond to the four remaining planes of the planning blackboard. Thus, the subject makes decisions about data--how long errands should take, how important individual errands are, what the consequences of a particular action might be, etc. He makes abstract planning decisions--what kinds of planning decisions might be useful. He makes meta-planning decisions--how to approach the problem and how to constrain and evaluate his plan. Finally, the subject makes executive decisions about how to allocate his cognitive resources during planning. While all of these decisions contribute to the planning process, it is difficult to see how one might incorporate them in a simple hierarchical planning structure.

Relative Advantages of Hierarchical versus Opportunistic Planning

We might also speculate on the relative merits of hierarchical versus opportunistic planning. The orderly, systematic nature of the top-down process and the simplicity of its hierarchical structure argue in its favor. The recent emphasis on "structured programming," a top-down approach to software engineering, reflects these merits (Cf., Dahl, Dykstra, & Hoare, 1972). One might also argue that top-down processes would minimize memory load (Cf., Thorndyke, 1978). The planner could restrict his attention to a single area of the hierarchy, rather than attending intermittently to several different areas of the planning space.

On the other hand, planning in tasks fraught with complexity and uncertainty might benefit from less of the discipline imposed by a top-down process. In such complex tasks, general, a priori solutions or problem-solving methods may not exist or may be computationally intractable. Even if some general approach were available, opportunistic planning would free the planner of the burden of maintaining a structurally integrated plan at each decision point. Instead, the planner could formulate and pursue promising partial plans as opportunity suggested.

More importantly, a multi-directional process might produce better plans. It certainly permits more varied plans than a top-down process does. If the planner always began with a fixed high-level plan, he could refine it into only a limited

number of different detailed plans. The bottom-up component in multi-directional processing represents an important source of innovation in planning. Low-level decisions and related observations can inspire novel higher-level plans. We observed this in the errand-planning protocol, for example, when the subject generated a high-level design based on observations and decisions made at the lower procedure level. Similarly, Feitelson and Stefik (1977) observed that their expert geneticist deliberately exploited the potential for innovation in bottom-up processing:

> Thus, not only is the planning process largely event driven but sometimes steps are taken somewhat outside the plan of the experiment to make a possibly interesting observation. This kind of behavior reflects the convenience of making certain interesting observations while the equipment is set up. Often this is done to verify the successful completion of an experimental step, but sometimes the observations seem to correspond more to fishing for interesting possibilities. (p. 31)

Resolving the Two Points of View

We have argued against the earlier view of planning as a systematic, hierarchical, top-down process and in favor of the proposed model of planning as an opportunistic, heterarchical, multi-directional process. We must consider, however, the possibility that both models have merit and simply apply to different situations. In particular, it appears plausible that

the top-down model could accurately describe planning by a practiced planner working on a familiar, constrained problem. In this kind of situation, the planner may have well-learned, reliable abstract plans for dealing with the problem. His extensive experience may support the application of standard methods for systematically refining his abstract plans.

A study by Byrne (1977) provides some support for this conjecture. His subjects planned dinner menus, a task with which they had considerable experience. As one might expect, Byrne's subjects appeared to plan menus by deciding on type of dinner (e.g., Chinese dinner, Christmas dinner), course (e.g., first course, main course, dessert), dish (e.g., roast beef, turkey), and accompaniments to a dish (e.g., cranberry sauce, mashed potatoes). This is a nice example of a hierarchical planning structure. In addition, Byrne's subjects appeared to make decisions within this structure in a top-down fashion.

On the other hand, subjects who performed the errand-planning task undoubtedly had considerable experience as well. Yet they did not exhibit systematic top-down planning behavior. We attribute their opportunistic planning activity to the greater complexity and difficulty of the task.

Additional support for the opportunistic model of expert planning derives from Feitelson's and Stefik's (1977) study of the experiment-planning behavior of an expert molecular geneticist. Part of that report follows:

The experiments described here reflect a combination of

> goal driven behavior and event driven behavior... If there were no goals, behavior might seem very erratic and follow no general course. If there is no event driven component to the planning process, then the experimental procedure must admit no feedback or changes of plans as a result of observations. Thus, no advantage will be made of fortunate observations. What is being suggested here is that the planning in this experiment involved far more exploitation of events and changes of plan according to the events than the authors had anticipated. (p. 30)

Thus, while the idea that "experts" at a planning task should be more likely to plan in a top-down fashion has some intuitive appeal, the available evidence is equivocal.

We can also attempt to resolve the apparent conflict between the two models by viewing the top-down model as one particular instantiation of the opportunistic model. Earlier, we discussed the importance of the problem-solving method a planner brings to bear on his task. This decision can have a major impact on subsequent executive decisions and, consequently, on the planner's progress through the remaining levels of the blackboard. In particular, a planner might adopt a "define and successively refine" problem-solving method. Given strict adherence to this method, the planner's formulation of decisions on the plan plane would indeed proceed in a systematic top-down fashion. These are exactly the decisions modeled in the earlier work on top-down planning.

Note that "define and refine" is only one of many problem-solving methods adoptable in the framework of the opportunistic model. Thus, the question is no longer which "model" is correct, but rather, under what circumstances do planners bring various problem-solving methods to bear? We have suggested familiarity with and complexity of the planning task as two potentially important factors. We need to investigate the effects of these and other factors. We should also ask which problem-solving methods work best for different kinds of problems. We have discussed some of the relative advantages of strictly top-down planning versus unbridled opportunistic planning. We need more research in this area as well.

Relation to Organizational Planning Models

Interestingly, organizational planning appears to mimic the proposed conception of individual planning. In organizational planning, various people make different kinds of decisions regarding different aspects of a plan. The various people correspond to the separate specialists postulated in the model of individual planning. The kinds of decisions these people make correspond to the kinds of decisions that appear on the planning blackboard. For example, Preston and Henning (1961) refer to several kinds of decisions:

> Objectives and goals are thought of as statements of the purposes for which our organized group has been formed... Procedures are more detailed and specific guides which are particularly helpful to operating personnel in the

> daily performance of their routine tasks. Policies are general guides to future decision-making that are intended to shape those decisions so as to maximize their contribution to the goals of the enterprise... [Methods are] less general ... [and] of less significance than policies. They are relatively detailed, complete, and specific plans for the guidance of business activities. (pp. 5-9)

Further, each of these kinds of decisions can change in the course of planning. For example, "objectives may change as conditions warrant" and "as a plan is prepared, it will often require the creation of new policies or the redefinition of existing policy" (Preston & Henning, 1961, pp. 4-6). Similarly, March (1972, p. 6) suggests: "Suppose we treat action as a way of creating interesting goals at the same time we treat goals as a way of suggesting actions."

In a retrospective analysis of the design of the highly successful PDP-11 computer, Bell (1977) describes a similarly organized planning process:

> Because of the many pressures on the design, the planning was asynchronous and diffuse; development was distributed throughout the company. This sort of decentralized design organization provides a system of checks and balances, but often at the expense of perfect hardware compatibility. This compatibility can hopefully be provided in the software and at lower cost to the user. (p. 9)

Thus, various specialists took responsibility for designing different components of the PDP-11. In addition, the specialists worked "asynchronously." That is, while they undoubtedly communicated with one another and influenced one another's decisions, the specialists did not coordinate their activities in any systematic way. While Bell identifies both costs and benefits associated with this kind of planning process, he also notes that the resulting system "exceeded the design goals" (p. 7).

Finally, March (1972) advances a position similar to ours regarding the value of opportunism (which he calls "playfulness") in planning and other decisionmaking:

> A second requirement...is some strategy for suspending rational imperatives toward consistency... A strict insistence on purpose, consistency, and rationality limits our ability to find new purposes. Play relaxes that insistence to allow us to act "unintelligently," or "irrationally," or "foolishly"--to explore alternative ideas of possible purposes and alternative concepts of behavioral consistency... For organizations and for individuals, reason and intelligence have had the unnecessary consequence of inhibiting the development of purpose into more complicated forms of consistency. In order to move away from that position, we need to find some way of helping individuals and organizations to

> experiment with doing things for which they have no good reason... We encourage organizations by permitting (and insisting on) some temporary relief from control, coordination, and communication... It preserves the virtues of consistency while stimulating change. (pp. 6-7)

GENERAL DISCUSSION

Status of the Proposed Model

The opportunistic model provides a nice description of the planning process. Its emphasis on the importance of both high-level goals and low-level details gives the model ecological validity. Obviously, goals exert a powerful effect on the planning process. However, in "real life," people frequently plan in the context of continually changing circumstances and capabilities. Effective planners must take these changing "details" into account as well.

The opportunistic model also predicts the gross characteristics of planning captured in thinking aloud protocols: that planning is opportunistic, incremental, multi-directional, and heterarchical. At a finer level of analysis, the model provides a comprehensive, if not complete, taxonomy of the kinds of decisions people apparently make while planning. Finally, the model provides a mechanism, the pattern-directed activity of independent specialists, whereby diverse decisions can influence one another. This notion of pattern-directed specialists

generalizes the view of production rules as simple, symbol-manipulating operations acting upon the contents of short-term memory (Newell & Simon, 1972). In this more general view, different specialists exploit arbitrarily complex reasoning procedures to generate or modify decisions (Cf., Hayes-Roth, Waterman, & Lenat, 1978, pp. 578-580).

The opportunistic model seems, at first glance, fairly complex. It postulates five different conceptual "planes" of decisions and several levels of abstraction within each of those planes. It postulates numerous planning specialists whose simultaneous efforts to participate in the planning process require the supervision of a fairly sophisticated executive. Although a number of comparably complex models have proved fruitful in the last few years (Cf., Anderson, 1976; Anderson & Bower, 1973; Rumelhart, Lindsay, & Norman, 1972; Winograd, 1972), most of us still adhere to the law of parsimony, preferring simpler models to complex models. Accordingly, we offer two reasons for advancing a model as complex as the opportunistic model.

First, planning is a complex process. It takes a considerable amount of time to formulate plans. Planners pass through numerous intermediate stages and consider a variety of information. They bring a variety of knowledge to bear on the problem. They perform many diverse sub-tasks in the course of planning. They evaluate numerous alternative plans and sub-plans and change their decisions repeatedly while planning. We feel that given the complexity of the planning process, the

opportunistic model is not unduly complex.

Second, the model describes the data well. People reach each of the several kinds of intermediate decisions postulated. They use each of these kinds of decisions in determining various subsequent decisions. They work intermittently on various aspects of a plan in progress. Simpler models, such as the top-down, hierarchical models discussed earlier, do not reflect this richness of the planning process.

In addition to being a very powerful model, the opportunistic model is vulnerable to data. In this paper, we discussed a thinking aloud protocol from the errand-planning task. The model predicted the gross characteristics of the protocol. Further, these characteristics appeared in most of the other twenty-nine protocols we collected. In addition, we can test the model's assumptions using conventional experimental methodology. For example, we can evaluate the psychological validity of the postulated organization and levels of the planning space. We can test the power of decisions at a given level of the planning space to influence subsequent decisions at other levels of the planning space. We can investigate the hypothesized impact of alternative problem-solving methods on the performance of particular planning tasks. Experiments along these lines are in progress.

Implications for Future Research

Control. As discussed in the first section of this

paper, planning is the first stage in a two-stage planning and control process. In this paper, we concentrated on the first stage, formulation of a satisfactory plan, and did not directly address the problem of controlling its execution. However, the proposed model provides a firm basis for approaching that problem.

The control process entails monitoring and guiding the execution of a plan to a successful conclusion. The planner must monitor progress toward the goal. If he determines that progress is unsatisfactory, he must alter the plan accordingly. The planner must also attend to unanticipated opportunities and obstructions. If a particularly attractive opportunity arises, the planner may decide to alter his plan in order to exploit that opportunity. If some obstruction appears, the planner must decide how to circumvent it.

The model provides a starting point for modeling control in its provisions for simulated execution of a plan in progress and plan revision based on the outcome of that simulation. We need only substitute actual execution for simulated execution to have a working model of control. Presumably, the planner works on the same blackboard during both planning and control. In both cases, the planner performs both step-wise and event-driven analyses, recording the results at the appropriate location on the blackboard (e.g., as attributes of the plan in progress, recorded on the plan plane, or attributes of the world, recorded in the knowledge base). Particular specialists respond to the occurrence of such information by suggesting appropriate

modifications to the plan.

Learning. The opportunistic model also provides a framework for studying learning. Many researchers have postulated largely bottom-up learning mechanisms (Cf., Franks & Bransford, 1971; B. Hayes-Roth, 1977; F. Hayes-Roth, 1978; Hayes-Roth & Hayes-Roth, 1977; Hayes-Roth & McDermott, 1978; Hebb, 1949; La Berge & Samuels, 1974; Mandler, 1962; Martin & Noreen, 1974; McGuire, 1961; Neumann, 1974; Posner, 1969; Posner & Keele, 1970; Reed, 1972; Reitman & Bower, 1973). These researchers assume that people learn simple "patterns" first. Given a set of simple patterns, people can combine patterns to form more complex patterns and abstract the distinguishing features of sets of related patterns. These assumptions imply that, with experience, people acquire successively more encompassing and more abstract knowledge. The opportunistic model explicates the categorical levels of planning knowledge that people presumably acquire. By adopting the assumption that people learn specific, low-level patterns before learning more complex, abstract patterns, the model can predict learning profiles.

Consider the knowledge exhibited in the errand-planning task. Under the preceding assumptions, people should learn to formulate low-level plans before learning to formulate high-level plans. For example, we would expect to observe children planning how to get from one point to another (operations) before they can plan more general sequences of point-to-point connections (procedures). Similarly, we would expect children to plan

sequences of point-to-point connections (procedures) before they can plan their travel behavior more generally, without regard to specific point-to-point connections (designs). Considering the relationships among planes, we would expect people to learn to plan what to do in the world (plan plane) before learning how to plan abstractly (plan-abstractions plane). We would expect people to learn how to form meta-plans and make intelligent executive decisions last of all (Cf., Soloway & Riseman, 1977).

Given the basic assumptions of the learning models cited above, the model can also predict the nature of the abstract, higher-level planning knowledge that people should acquire from particular prior planning experiences. For example, suppose an individual made good use of the "go to the closest errand next" strategy. The individual might have several experiences in which successive applications of this strategy lead to implicit spatial clustering of the planned errands. This could provide a basis for inducing a more general "clustering" scheme (Cf., Fikes, Hart & Nilsson, 1972; F. Hayes-Roth, 1977).

Expertise. These considerations lead naturally to the question of defining expertise. One of the most valuable results we could obtain would be an understanding of why some people plan well while others plan poorly. What makes an expert planner? The general issue of expertise has received considerable attention, and some progress has been made toward characterizing expertise in domains other than planning (Bhaskar & Simon, 1977; Chase & Simon, 1973; Marples, 1974; Simon & Simon, 1977; Reitman, 1976). The proposed planning model provides a rich framework for

modeling expert planning, suggesting several sources of expertise as discussed below.

Expert planners probably use more levels of the planning space. In accordance with the predicted order of acquisition discussed above, experts probably use more abstract levels within a plane and use more abstract planes in general. It seems especially likely that experts would do more meta-planning and make more sophisticated executive decisions than non-experts.

Previous studies of expert problem-solving behavior provide some support for the latter conjecture. For example, Thorndyke and Stasz (personal communication) find that expert map learners systematically sample and rehearse the various regions of a target map and intentionally focus on just those areas of the map they have not yet mastered. While sampling strategies vary among experts, all experts exhibit clearly defined strategies. Several other studies (Bhaskar & Simon, 1977; Marples, 1974; Simon & Simon, 1977) suggest that the expert's experience with a problem domain enables him to distinguish between problems which he can solve simply by working bottom-up from the initial conditions and those which require more complex means-ends analysis. Thus, experts apply bottom-up strategies to some problems and means-ends analysis to others. These results suggest that expertise requires more than an ability to apply one "correct" or "optimal" executive control structure. Rather, expertise requires an ability to apply alternative executive control structures as appropriate for a given problem.

Expert planners also probably employ more specialists. Given our hypothesis that they operate on more abstract levels of the planning space, we expect experts to have certain top-down specialists that non-experts do not have. In addition, expert planners probably have more powerful specialists. For example, they might have certain bottom-up specialists that produce valuable innovations in their plans.

Expert planners probably excel in bringing knowledge to bear on a developing plan. For example, they probably have more accurate or more complete knowledge bases. These factors would produce domain-specific planning experts rather than general planning experts. Studies of chess and Go (Chase & Simon, 1973; Reitman, 1976), showing that experts recognize more complex board configurations than novices, support this conjecture. Experts might also have more effective specialists for performing certain computations on information in the knowledge base. Following similar reasoning, Simon and Simon (1977) note the importance of an expert physicist's superior algebraic and arithmetic skill.

Expert planners might also simulate execution of tentative plans more accurately. Simon and Simon (1977) provide some support for this conjecture. They suggest that an expert physicist can represent a problem internally in terms of physical relationships among objects. He then simulates the behavior of these objects, as specified in the problem, to aid his other problem-solving activity.

Finally, expert planners might excel in certain basic

cognitive functions (Cf., Hunt, 1978). For example, expert planners might have superior memory. One of the apparent disadvantages of the opportunistic approach to planning is that it places a burden on working memory. Superior memory capacity would lighten this burden for experts. In addition, expert planners might have more flexible attentional processes, enabling them to exploit the distributed activities assumed under the proposed model (Cf., Gopher & Kahneman, 1971; Kahneman, Ben-Ishai, & Lotan, 1973; Keele, Neill, & deLemos, 1978).

Conclusions

In sum, the opportunistic model provides a comprehensive framework for modeling planning and associated cognitive processes. The model draws on earlier theoretical work from cognitive, organizational, and computational domains. It reinterprets the strongest points of the earlier models and combines them with its own assumptions regarding multi-directionality and opportunism in a heterarchical plan structure. The resulting model represents a qualitatively new approach to planning. It is a powerful model, flexible in the face of a complex cognitive behavior, yet vulnerable to data. Finally, the model provides a rich framework for studying control of plan execution, the acquisition of planning skills and the development of expertise.

REFERENCES

Aho, A. V., Hopcroft, J. E., & Ullman, J. D. The design and analysis of computer algorithms. Reading, Massachusetts: Addison-Wesley, 1974.

Anderson, J. R. Language, memory, and thought. Hillsdale, N.J.: Lawrence Erlbaum, 1976.

Anderson, J. R., & Bower, G. H. Human associative memory. Washington, D.C.: V. H. Winston, 1973.

Atwood, M. E., & Polson, P. G. A process model for water jug problems. Cognitive Psychology, 1976, 8, 191-216.

Bhaskar, R., & Simon, H. A. Problem solving in semantically rich domains: An example from engineering thermodynamics. Cognitive Science, 1977, 1, 192-215.

Bell, C. G. What have we learned from the PDP-11? In A. K. Jones (Ed.), Perspectives on computer science. New York: Academic Press, 1977.

Byrne, R. Planning meals: Problem-solving on a real data-base. Cognition, 1977, 5, 287-332.

CMU Computer Science Research Group. Summary of the CMU five-year ARPA effort in speech understanding research. Technical Report, Carnegie-Mellon University, 1977.

Chase, W. G., & Simon, H. A. Perception in chess. Cognitive Psychology, 1973, 4, 55-81.

Christophides, N. Graph theory: An algorithmic approach. New York: Academic Press, 1975.

Dahl, O. J., Dykstra, E. W., & Hoare, C. A. R. Structured programming. New York: Academic Press, 1972.

Engelmore, R. S., & Nii, H. P. A knowledge-based system for the interpretation of protein x-ray crystallographic data. Report No. STAN-CS-77-589, Stanford University, 1977.

Erman, L. D., & Lesser, V. R. 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., & Newell, A. 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.

Feitelson, J., & Stefik, M. A case study of the reasoning in a genetics experiment. Heuristic Programming Project, Working Paper 77-18, Department of Computer Science, Stanford University, April 1977.

Fikes, R. E. Knowledge representation in automatic planning systems. In A. K. Jones (Ed.), Perspectives on computer science. New York: Academic Press, 1977.

Fikes, R. E., & Nilsson, N. J. STRIPS: A new approach to the

application of theorem proving to problem solving. Artificial Intelligence, 1971, 2, 189-203.

Fikes, R. E, Hart, P. E., & Nilsson, N. J. Learning and executing generalized robot plans. Artificial Intelligence, 1972, 3, 251-288.

Franks, J. J., & Bransford, J. D. Abstraction of visual patterns. Journal of Experimental Psychology, 1971, 90, 65-74.

Gopher, D, & Kahneman, D. Individual differences in attention and the prediction of flight criteria. Perceptual and Motor Skills, 1971, 33, 1335-1342.

Greeno, J. G. Hobbits and Orcs: Acquisition of a sequential concept. Cognitive Psychology, 1974, 6, 270-292.

Hayes-Roth, B. Evolution of cognitive structures and processes. Psychological Review, 1977, 84, 260-278.

Hayes-Roth, B., & Hayes-Roth, F. Concept learning and the recognition and classification of exemplars. Journal of Verbal Learning and Verbal Behavior, 1977, 16, 321-338.

Hayes-Roth, F. Learning by example. In A. Lesgold, J. Pellegrino, S. Fokkema, and R. Glaser (Eds.), Cognitive psychology and instruction. New York: Plenum, 1978.

Hayes-Roth, F., & Lesser, V. R. 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.

Hayes-Roth, F., & McDermott, J. An interference matching technique for inducing abstractions. Communications of the ACM, 1978, 21, 401-410.

Hayes-Roth, F., Waterman, D. A., & Lenat, D. B. Principles of pattern-directed inference systems. In D. A. Waterman and F. Hayes-Roth (Eds.), Pattern-directed inference systems. New York: Academic Press, 1978.

Hebb, D. O. Organization of behavior. New York: Wiley, 1949.

Hunt, E. Mechanics of verbal ability. Psychological Review, 1978, 85, 109-130.

Jeffries, R., Polson, P., Razran, L., & Atwood M. E. A process model for missionaries-cannibals and other river-crossing problems. Cognitive Psychology, 1977, 9, 412-440.

Kahneman, D., Ben-Ishai, R., & Lotan, M. Relation of a test of attention to road accidents. Journal of Applied Psychology, 1973, 58, 113-115.

Keele, S. W., Neill, W. T., de Lemos, S. M. Individual differences in attentional flexibility. Center for Cognitive and Perceptual Research Technical Report, 1978.

LaBerge, D., & Samuels, S. J. Toward a theory of automatic information processing in reading. Cognitive Psychology, 1974, 6, 293-323.

Lesser, V. R., Fennell, R. D., Erman, L. D., & Reddy, D. R. Organization of the Hearsay-II speech understanding system.

IEEE Transactions on Acoustics, Speech and Signal Processing, ASSP-23, (Jan. 1975), 11-23.

Lesser, V. R., & Erman, L. D. A retrospective view of the Hearsay-II architecture. Proceedings of the Fifth International Joint Conference on Artificial Intelligence, Boston, Mass., 1977, 790-800.

Mandler, G. From association to structure. Psychological Review, 1962, 69, 415-427.

March, J. The technology of foolishness. Stanford Alumni Almanac Vol. II, October 1972.

Marples, D. L. Argument and technique in the solution of problems in mechanics and electricity, Department of Engineering, University of Cambridge, England, 1974.

Martin, E., & Noreen, D. L. Serial learning: Identification of subjective subsequences. Cognitive Psychology, 1974, 6, 421-435.

McGuire, W. J. A multipurpose model for paired-associate learning. Journal of Experimental Psychology, 1961, 62, 335-347.

Miller, G. A., Galanter, E., & Pribram, D. H. Plans and the structure of behavior. New York: Holt, Rinehart and Winston, Inc., 1960.

Neumann, P. G. An attribute frequency model for the abstraction of prototypes. Memory and Cognition, 1974, 2, 241-248.

Newell, A. You can't play 20 questions with nature and win: Projective comments on the papers of this symposium. In W. G. Chase (Ed.), Visual information processing. New York: Academic Press, 1973.

Newell, A., Shaw, J. C., & Simon, H. A. Chess playing programs and the problem of complexity. IBM Journal of Research and Development, 1958, 320-335. (Reprinted in E. A. Feigenbaum & J. Feldman (Eds.), Computers and thought. New York: McGraw-Hill, 1963.)

Newell, A., & Simon, H. A. Human problem solving. Engelwood Cliffs, N.J.: Prentice-Hall, Inc., 1972.

Nii, H. P., & Feigenbaum, E. A. Rule-based understanding of signals. In D. A. Waterman and F. Hayes-Roth (Eds.), Pattern-directed inference systems. New York: Academic Press, 1978.

Posner, M. I. Abstraction and the process of recognition. In G. H. Bower and J. T. Spence (Eds.), Psychology of learning and motivation. New York: Academic Press, 1969, Vol. 3.

Posner, M. I., & Keele, S. W. Retention of abstract ideas. Journal of Experimental Psychology, 1970, 83, 304-308.

Prager, J., Nagin, P., Kohler, R., Hanson, A., & Riseman, E. Segmentation processes in the VISIONS system. Proceedings of the Fifth International Joint Conference on Artificial Intelligence, Boston, Mass., 1977.

Preston, P., & Henning, D. A. Planning theory. Englewood Cliffs, N.J.: Prentice-Hall, 1961.

Reed, S. K. Pattern recognition and categorization. Cognitive Psychology, 1972, 3, 382-407.

Reed, S. K., Ernst, G. W., & Banerji, R. The role of analogy in transfer between similar problem states. Cognitive Psychology, 1974, 6, 436-450.

Reitman, J. S. Skilled perception in Go: Deducing memory structures from inter-response times. Cognitive Psychology, 1976, 8, 336-356.

Reitman, J. S., & Bower, G. H. Storage and later recognition of exemplars of concepts. Cognitive Psychology, 1973, 4, 194-206.

Rumelhart, D. E. Toward an interactive model of reading. Technical Report 56, Center for Human Information Processing, University of California at San Diego, La Jolla, Ca., 1976.

Rumelhart, D., Lindsay, P. H., & Norman, D. A. A process model for long-term memory. In E. Tulving & W. Donaldson (Eds.), Organization of memory. New York: Academic Press, 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.

Selfridge, O. Pandemonium: A paradigm for learning. Symposium on the mechanization of thought. London: HM Stationery Office, 1959.

Simon, D. P., & Simon, H. A. Individual differences in solving physics problems. C.I.P. Working Paper No. 342, Department of Psychology, Carnegie-Mellon University, 1977.

Soloway, E. M., & Riseman, E. M. Knowledge-directed learning. Proceedings of the Workshop on Pattern-Directed Inference Systems, special edition of the SIGART Newsletter. New York: Association for Computing Machinery, June, 1977.

Sussman, G. J. A computational model of skill acquisition. AI TR-297, Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Mass., 1973.

Thomas, J. C., Jr. An analysis of behavior in the Hobbits-Orcs problem. Cognitive Psychology, 1974, 6, 257-269.

Thorndyke, P. W. Pattern-directed processing of knowledge from text. In D. A. Waterman and F. Hayes-Roth (Eds.), Pattern-directed inference systems. New York: Academic Press, 1978.

Winograd, T. A program for understanding natural language. Cognitive Psychology, 1972, 3.

FOOTNOTES

1. ONR Contract N00014-78-C-0039, NR 157-411 supported this research. We thank Bob Anderson, Ed Feigenbaum, Penny Nii, Perry Thorndyke and members of the Rand Cognitive Sciences Brownbag for many valuable discussions of this research. Bob Anderson, Bill Faught, Phil Klahr, Stan Rosenschein, and Bob Wesson provided useful comments on an earlier version of this manuscript. Send reprint requests to: Barbara Hayes-Roth, The Rand Corporation, 1700 Main Street, Santa Monica, California, 10406.

Navy

1 Dr. Jack Adams
Office of Naval Research Branch
223 Old Marylebone Road
London, NW, 15th England

1 Dr. Ed Aiken
Navy Personnel R&D Center
San Diego, CA 92152

1 Dr. Jack R. Borsting
Provost & Academic Dean
U.S. Naval Postgraduate School
Monterey, CA 93940

1 Dr. Robert Breaux
Code N-71
NAVTRAEQUIPCEN
Orlando, FL 32813

1 Dept. of the Navy
CHNAVMAT (NMAT 034D)
Washington, DC 20350

1 Chief of Naval Education and
Training Support)-(01A)
Pensacola, FL 32509

1 Dr. Charles E. Davis
ONR Branch Office
536 S. Clark Street
Chicago, IL 60605

1 Mr. James S. Duva
Chief, Human Factors Laboratory
Naval Training Equipment Center
(Code N-215)
Orlando, FL 32813

5 Dr. Marshall J. Farr, Director
Personnel & Training Research
Programs
Office of Naval Research
(Code 458)
Arlington, VA 22217

1 Dr. Pat Federico
Navy Personnel R&D Center
San Diego, CA 92152

Navy

1 CDR John Ferguson, MSC, USN
Naval Medical R&D Command
(Code 44)
National Naval Medical Center
Bethesda, MD 20014

1 Dr. John Ford
Navy Personnel R&D Center
San Diego, CA 92152

1 Dr. Eugene E. Gloye
ONR Branch Office
1030 East Green Street
Pasadena, CA 91101

1 Capt. D.M. Gragg, MC, USN
Head, Section on Medical Education
Uniformed Services Univ. of the
Health Sciences
6917 Arlington Road
Bethesda, MD 20014

1 Dr. Steve Harris
Code L522
NAMRL
Pensacola, FL 32508

1 CDR Robert S. Kennedy
Naval Aerospace Medical and
Research Lab
Box 29407
New Orleans, LA 70189

1 Dr. Norman J. Kerr
Chief of Naval Technical Training
Naval Air Station Memphis (75)
Millington, TN 38054

1 Chairman, Leadership & Law Dept.
Div. of Professional Development
U.S. Naval Academy
Annapolis, MD 21402

1 Dr. James Lester
ONR Branch Office
495 Summer Street
Boston, MA 02210

Navy

1 Dr. William L. Maloy
Principal Civilian Advisory for Education and Training
Naval Training Command, Code 00A
Pensacola, FL 32508

1 Dr. James McBride
Code 301
Naval Personnel R&D Center
San Diego, CA 92152

2 Dr. James McGrath
Navy Personnel R&D Center
Code 306
San Diego, CA 92152

1 Dr. William Montague
Navy Personnel R&D Center
San Diego, CA 92152

1 Commanding Officer
U.S. Naval Amphibious School
Coronado, CA 92155

1 Commanding Officer
Naval Health Research Center
Attn: Library
San Diego, CA 92152

1 CDR Paul Nelson
Naval Medical R&D Command
Code 44
National Naval Medical Center
Bethseda, MD 20014

1 Library
Navy Personnel R&D Center
San Diego, CA 92152

6 Commanding Officer
Naval Research Laboratory
Code 2627
Washington, DC 20390

1 John Olsen
Chief of Naval Education and Training Support
Pensacola, FL 32509

1 Office of Naval Research
Code 200
Arlington, VA 22217

Navy

1 Office of Naval Research
Code 437
800 N. Quincy Street
Arlington, VA 22217

1 Scientific Director
Office of Naval Research
Scientific Liaison Group/Tokyo
American Embassy
APO San Francisco, CA 96503

1 Scientific Advisor to the Chief of Naval Personnel
Naval Bureau of Personnel (PERS OR)
Rm. 4410, Arlington Annex
Washington, DC 20370

1 Dr. Richard A. Pollak
Academic Computing Center
U.S. Naval Academy
Annapolis, MD 21402

1 Mr. Arnold I. Rubinstein
Human Resources Program Manager
Naval Material Command (0344)
Rm. 1044, Crystal Plaza #5
Washington, DC 20360

1 Dr. Worth Scanland
Chief of Naval Education and Training
Code N-5
NAS, Pensacola, FL 32508

1 A.A. Sjoholm
Tech. Support, Code 201
Navy Personnel R&D Center
San Diego, CA 92152

1 Mr. Robert Smith
Office of Chief of Naval Operations
OP-987E
Washington, DC 20350

1 Dr. Alfred F. Smode
Training Analysis & Evaluation Group (TAEG)
Dept. of the Navy
Orlando, FL 32813

1 CDR Charles J. Theisen, Jr., MSC, USN
Head Human Factors Engineering Div.
Naval Air Development Center
Warminster, PA 18974

Navy

1 W. Gary Thomson
Naval Ocean Systems Center
Code 7132
San Diego, CA 92152

Army

1 ARI Field Unit-Leavenworth
P.O. Box 3122
Ft. Leavenworth, KS 66027

1 HQ USAREUE & 7th Army
ODCSOPS
USAAREUE Director of GED
APO New York 09403

1 Commandant
U.S. Army Infantry School
Ft. Benning, GA 31905
Attn: ATSH-I-V-IT (Cpt. Hinton)

1 Dr. James Baker
U.S. Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

1 Dr. Ralph Dusek
U.S. Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

1 Dr. Frank J. Harris
U.S. Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

1 Dr. A. Hyman
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

1 Dr. Ed Johnson
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

1 Dr. Milton S. Katz
Individual Training & Skill
Evaluation Technical Area
U.S. Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

1 Dr. Harold F. O'Neil, Jr.
Attn: PERI-OK
5001 Eisenhower Avenue
Alexandria, VA 22333

Army

1 Director, Training Development
U.S. Army Administration Center
Attn: Dr. Sherrill
Ft. Benjamin Harrison, IN 46218

1 Dr. Joseph Ward
U.S. Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Air Force

1 Dr. G.A. Eckstrand
AFHRL/AS
Wright-Patterson AFB, OH 45433

1 Dr. Alfred R. Fregly
AFOSR/NL, Bldg. 410
Bolling AFB, DC 20332

1 CDR Mercer
CNET Liaison Officer
AFHRL/Flying Training Div.
Williams AFB, AZ 85224

1 Research Branch
AFMPC/DPMYP
Randolph AFB, TX 78148

1 Dr. Marty Rockway (AFHRL/TT)
Lowry AFB
Colorado 80230

1 Brian K. Waters, Maj., USAF
Chief, Instructional Tech. Branch
AFHRL
Lowry AFB, CO 80230

Marines

1 Director, Office of Manpower Utilization
HQ, Marine Corps (MPU)
BCB, Bldg. 2009
Quantico, VA 22134

1 Dr. A.L. Slafkosky
Scientific Advisor (Code RD-1)
HQ, U.S. Marine Corps
Washington, DC 20380

Coast Guard

1 Mr. Joseph J. Cowan, Chief
Psychological Research (G-P-1/62)
U.S. Coast Guard HQ
Washington, DC 20590

Other DoD

1 Dr. Stephen Andriole
Advanced Research Projects Agency
1400 Wilson Blvd.
Arlington, VA 22209

12 Defense Documentation Center
Cameron Station, Bldg. 5
Alexandra, VA 22314
Attn: TC

1 Dr. Dexter Fletcher
Advanced Research Projects Agency
1400 Wilson Blvd.
Arlington, VA 22209

1 Military Assistant for Human Resources
Office of the Director of Defense Research & Engineering
Rm. 3D129, The Pentagon
Washington, DC 20301

1 Director, Research & Data
OSD/MRA&L (Rm. 3B919)
The Pentagon
Washington, DC 20301

Other Govt

1 Dr. Susan Chipman
Basic Skills Program
National Institute of Education
1200 19th Street, NW
Washington, DC 20208

1 Dr. Joseph Markowitz
Office of Research and Development
Central Intelligence Agency
Washington, DC 20205

1 Dr. Thomas G. Sticht
Basic Skills Program
National Institute of Education
1200 19th Street, NW
Washington, DC 20208

1 Dr. Joseph L. Young, Director
Memory & Cognitive Processes
National Science Foundation
Washington, DC 20550

Non Govt

1 Prof. Earl A. Alluisi
Dept. of Psychology
Code 287
Old Dominion University
Norfolk, VA 23508

1 Dr. John R. Anderson
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

1 Dr. Michael Atwood
Science Applications Institute
40 Denver Tech. Center West
7935 E. Prentice Avenue
Englewood, CO 80110

1 1 Psychological Research Unit
Dept. of Defense (Army Office)
Campbell Park Offices
Canberra ACT 2600, Australia

1 Dr. Nicholas A. Bond
Dept. of Psychology
Sacramento State College
600 Jay Street
Sacramento, CA 95819

1 Dr. Lyle Bourne
Department of Psychology
University of Colorado
Boulder, CO 80302

1 Dr. John Seeley Brown
Bolt Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02138

1 Dr. John B. Carroll
Psychometric Lab
Univ. of No. Carolina
Davie Hall 013A
Chapel Hill, NC 27514

1 Dr. William Chase
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

Non Govt

1 Dr. Micheline Chi
Learning R&D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

1 Dr. Allan M. Collins
Bolt Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02138

1 Dr. John J. Collins
Essex Corporation
201 N. Fairfax Street
Alexandria, VA 22314

1 Dr. Meredith Crawford
5605 Montgomery Street
Chevy Chase, MD 20015

1 Dr. Ruth Day
Center for Advanced Study
in Behavioral Sciences
202 Junipero Serra Blvd.
Stanford, CA 94305

1 Dr. Hubert Dreyfus
Department of Psychology
University of California
Berkeley, CA 94720

1 Major I.N. Evonic
Canadian Forces Pers. Applied
Research
1107 Avenue Road
Toronto, Ontario, Canada

1 Dr. Ed Feigenbaum
Department of Computer Science
Stanford University
Stanford, CA 94305

1 Dr. Victor Fields
Dept. of Psychology
Montgomery College
Rockville, MD 20850

1 Dr. Edwin A. Fleishman
Advanced Research Resources Organ.
8555 Sixteenth Street
Silver Spring, MD 20910

Non Govt

1 Dr. John R. Frederiksen
Bolt Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02138

1 Dr. Robert Glaser
LRDC
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

1 Dr. Ira Goldstein
Xerox Palo Alto Research Center
3333 Coyote Road
Palo Alto, CA 04304

1 Dr. James G. Greeno
LRDC
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

1 Dr. Ron Hambleton
School of Education
University of Massachusetts
Amherst, MA 01002

1 Dr. Richard S. Hatch
Decision Systems Assoc., Inc.
350 Fortune Terrace
Rockville, MD 20854

1 Dr. James R. Hoffman
Department of Psychology
University of Delaware
Newark, DE 19711

1 Library
HumRRO/Western Division
27857 Berwick Drive
Carmel, CA 93921

1 Dr. Earl Hunt
Dept. of Psychology
University of Washington
Seattle, WA 98105

1 Mr. Gary Irving
Data Sciences Division
Technology Services Corporation
2811 Wilshire Blvd.
Santa Monica, CA 90403

Non Govt

1 Dr. Lawrence B. Johnson
Lawrence Johnson & Assoc., Inc.
Suite 502
2001 S Street, NW
Washington, DC 20009

1 Dr. Arnold F. Kanarick
Honeywell, Inc.
2600 Ridgeway Pkwy
Minneapolis, MN 55413

1 Dr. Steven W. Keele
Dept. of Psychology
University of Oregon
Eugene, OR 97403

1 Dr. Walter Kintsch
Department of Psychology
University of Colorado
Boulder, CO 80302

1 Dr. David Kieras
Department of Psychology
University of Arizona
Tucson, AZ 85721

1 Mr. Marlin Kroger
1117 Via Goleta
Palos Verdes Estates, CA 90274

1 LCOL C.R.J. Lafleur
Personnel Applied Research
National Defense HQS
101 Colonel By Drive
Ottawa, Canada K1A 0K2

1 Dr. Jill Larkin
SESAME
c/o Physics Department
University of California
Berkeley, CA 94720

1 Dr. Robert R. Mackie
Human Factors Research, Inc.
6780 Cortona Drive
Santa Barbara Research Pk.
Goleta, CA 93017

1 Mr. Mark Miller
Massachusetts Institute of
Technology
Artificial Intelligence Lab
545 Tech Square
Cambridge, MA 02139

Non Govt

1 Dr. Richard B. Millward
Dept. of Psychology
Hunter Lab.
Brown University
Providence, RI 82912

1 Dr. Donald A. Norman
Dept. of Psychology C-009
Univ. of California, San Diego
La Jolla, CA 92093

1 Dr. Jesse Orlansky
Institute for Defense Analysis
400 Army Navy Drive
Arlington, VA 22202

1 Dr. Seymour A. Papert
Massachusetts Institute of
Technology
Artificial Intelligence Lab
545 Tech Square
Cambridge, MA 02139

1 Mr. Luigi Petrullo
2431 N. Edgewood Street
Arlington, VA 22207

1 Dr. Peter Polson
Dept. of Psychology
University of Colorado
Boulder, CO 80302

1 Dr. Diane M. Ramsey-Klee
R-K Research & System Design
3947 Ridgemont Drive
Malibu, CA 90265

1 Dr. Peter B. Read
Social Science Research Council
605 Third Avenue
New York, NY 10016

1 Dr. Mark D. Reckase
Educational Psychology Dept.
University of Missouri-Columbia
12 Hill Hall
Columbia, MO 65201

1 Dr. Fred Reif
SESAME
c/o Physics Department
University of California
Berkeley, CA 94720

Non Govt

1 Dr. Joseph W. Rigney
Univ. of So. California
Behavioral Technology Labs
3717 South Hope Street
Los Angeles, CA 90007

1 Dr. Andrew M. Rose
American Institutes for Research
1055 Thomas Jefferson St., NW
Washington, DC 20007

1 Dr. Ernst Z. Rothkopf
Bell Laboratories
600 Mountain Avenue
Murray Hill, NJ 07974

1 Prof. Fumiko Samejima
Dept. of Psychology
University of Tennessee
Knoxville, TN 37916

1 Dr. Walter Schneider
Dept. of Psychology
University of Illinois
Champaign, IL 61820

1 Dr. Allen Schoenfeld
SESAME
c/o Physics Department
University of California
Berkeley, CA 94720

1 Dr. Robert Singer, Director
Motor Learning Research Lab
Florida State University
212 Montgomery Gym
Tallahassee, FL 32306

1 Dr. Richard Snow
School of Education
Stanford University
Stanford, CA 94305

1 Dr. Robert Sternberg
Dept. of Psychology
Yale University
Box 11A, Yale Station
New Haven, CT 06520

1 Dr. Albert Stevens
Bolt Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02138

Non Govt

1 Dr. Patrick Suppes
Institute for Mathematical
Studies in the Social Sciences
Stanford University
Stanford, CA 94305

1 Dr. Kikumi Tatsuoka
Computer Based Education
Research Laboratory
252 Engineering Research Laboratory
University of Illinois
Urbana, IL 61801

1 Dr. Benton J. Underwood
Dept. of Psychology
Northwestern University
Evanston, IL 60201

1 Dr. Thomas Wallsten
Psychometric Laboratory
Davie Hall 013A
University of North Carolina
Chapel Hill, NC 27514

1 Dr. Claire E. Weinstein
Educational Psychology Dept.
Univ. of Texas at Austin
Austin, TX 78712

1 Dr. David J. Weiss
N660 Elliott Hall
University of Minnesota
75 E. River Road
Minneapolis, MN 55455

1 Dr. Susan E. Whitely
Psychology Department
University of Kansas
Lawrence, KS 66044