Seifert, C. M., Patalano, A. L., Hammond, K. J., & Converse, T. M. (1997). Experience and expertise: The role of memory in planning for opportunities. In P. J. Feltovich, K. M. Ford & R. R. Hoffman (Eds.), *Expertise in context* (pp. 101-123). Menlo Park, CA: AAAI Press/ MIT Press.

CHAPTER 4

Experience and Expertise: The Role of Memory in Planning for Opportunities

Colleen M. Seifert, Andrea L. Patalano, Kristian J. Hammond, & Timothy M. Converse

Introduction

A commonsense notion about expertise is that experts differ from novices due to the number of experiences they have had within a particular domain. The experience acquired over time may form the basis of the complex knowledge organization that distinguishes experts from novices. But in addition to measures such as "years on the job," the quality of specific experiences may play a critical role in expertise. The content of these individual episodes may be preserved in memory and serve as a continuing resource for experts when faced with later problems. As such, they may serve as a context of experience with which to evaluate and refine the abstract knowledge comprising expertise.

In this chapter, we discuss an approach to expertise that emphasizes the role of specific past experiences in memory. We examine expertise in planning tasks, where decisions are made about which goals to pursue, which plans to select, and when to execute them. These decisions are likely to benefit from considering past experiences with similar cases. Our approach emphasizes the notion that "reasoning is remembering"—remembering successful plans so they can be repeated, remembering failed plans so they can be avoided, and remembering pending plans so they can be attempted later. Our main hypothesis is that expert planners are able to anticipate the circumstances related to satisfying goals, and therefore maximize their ability to respond to opportunities when they arise. We examine this opportunistic planning capability as a touchstone feature of human expertise at its best.

In the sections of this chapter, we briefly review the role of experience in expertise, and the central role of past experiences within our framework, case-based planning. Then, we present the problem of *opportunism*—taking advantage of unanticipated circumstances to satisfy one's goals—as a challenge for any model of expertise. After a discussion of competing models of opportunistic planning, we define our own model and identify its cognitive mechanisms. We then evaluate this model through two research methodologies: computational modeling and experiments with human participants. By combining evidence from these two approaches, we hope to develop a more complete theory of how opportunistic planning is accomplished. Our account of the role of memory for specific past experiences in identifying later opportunities provides evidence for the importance of the context of specific experiences in expertise.

The Need for Experience-Based Expertise

Most classical AI work on planning includes a separation of planning from execution. These theories (e.g., Sussman 1975) emphasize complete preplanning in order to guarantee that only a "correct" plan will ever be executed. They accomplish this by deriving a complete plan, projecting its effects on a model of the world, and determining that its outcome is successful. But complete planners, such as Sussman's STRIPS model (1975), also require assuming that the planner has complete and flawless knowledge of the world, that the world is stable and unchanging, and that the planner has all the time it needs to plan before any action is required. These assumptions may not hold in many situations; instead, we find our knowledge is incomplete, the world constantly changes, and unanticipated objects and events are perceived. As we begin to work with more realistic domain models, new issues arise that make complete preplanning less feasible (Hammond, Marks, and Converse 1985):

The Costly Information Issue: A planner cannot count on having complete information about a domain. Understanding the world "correctly," storing new information in a usable form, and deducing all potential outcomes of information can be expensive if not impossible operations.

The Planning Complexity Issue: If a planner searches for a correct and safe plan by projecting forward the effects of initial steps, the computational complexity can quickly become too time-consuming to be practical.

The Execution-Time Issue: A plan which seemed correct when it was constructed may turn out to fail during execution. To cope, a planner must have some facility for replanning, recovery, and repair.

The Multiple Goals Issue: As a planner interacts with the world, it is sometimes confronted with multiple, simultaneous goals, that interact in unexpected ways.

The Postponed Execution Issue: Because more goals can be created than can be pursued simultaneously, it is often the case that some goals will have to be postponed, and returned to when the time is right to achieve them.

We argue that many real-world domains have these characteristics (e.g., military exercises, stock purchasing, news programming), and require an expert to function with uncertain information about the world. However, one paradigm that explicitly attempts to address these problems is case-based reasoning (Bareiss 1989; Carbonell 1981; Hammond 1989; Kolodner, Simpson, and Sycara 1985; Schank 1982). The case-based approach to planning may be a tractable alternative to complete planners because, like human planners, it does not depend on having a complete solution in advance.

Case-Based Planning

Case-based planning (a specialization of case-based reasoning for planning tasks) proposes that the way to deal with the combinatorics of planning and projection is to let past experiences tell the planner when and where plans work and don't work. Rather than preplanning, a planner should reuse past successful plans. Rather than projecting the effects of actions into the future, a planner should recall what the outcomes were in the past. Rather than simulating a plan to identify problematic interactions, a planner should recall and then avoid those that have cropped up before.

The core notion of case-based planning is that expertise depends on the development of domain knowledge (a "case library" containing both memories of specific experiences and generalizations, abstractions, and rules based on those experiences). By experiencing failures and successes within a domain, the expert planner is prepared to deal with future failures and opportunities when they occur. Planning from past experiences is supported by some psychological evidence for making use of past experiences in planning for new goals (Byrne 1977).

In case-based planning, the case library is built up incrementally through experience. The framework (Hammond 1989) includes seven case-based processes:

1. An ANTICIPATOR that predicts execution problems on the basis of the failures that have been seen by the planner in the past.
2. A RETRIEVER that searches memory for a plan that satisfies current goals while avoiding any problems that the ANTICIPATOR has predicted.
3. A MODIFIER that alters the plan found by the RETRIEVER to achieve any other goals from the input that are not yet satisfied.
4. A PROJECTOR that uses past cases indexed by plans to predict outcomes.
5. A STORER that places new plans in memory, indexed by the goals that they satisfy and the problems that they avoid.

6. A REPAIRER that operates when a plan fails and attempts to explain and correct problems encountered.
7. An ASSIGNER that uses the causal explanation built during repair to determine the features which will predict this failure in the future. This knowledge is used to index the failure for later access by the ANTICIPATOR.

These seven modules define the case-based planning approach, and are presented in depth by Hammond (1989). Together, these modules act to anticipate and avoid execution problems before they occur, apply a variety of repairs to faulty plans that it creates, and store plans in memory so that they can be used again. The case-based approach includes the power of rule-based reasoning, since generalizations can be built from commonalities among cases. However, the case library is also preserved in memory, and can be accessed as a resource for future reasoning.

One example of a case-based planner is CHEF, which operates in the domain of Szechwan cooking (Hammond 1989). Its input is a set of goals for different tastes, textures, ingredients, and types of dishes, and its output is a single recipe that satisfies multiple goals. It finds a past plan that satisfies as many of the most important goals as possible, and then modifies that plan to satisfy additional goals as well. Before searching memory for a plan to modify, CHEF examines the goals and recalls any failures that have occurred in the interactions among plans. If a related failure is found in memory, CHEF adds a goal of avoiding the failure to its set of goals to satisfy, and this new goal is also used to search for a plan.

For example, suppose CHEF anticipates that stir frying chicken with snow peas can lead to soggy snow peas (because the chicken will sweat liquid into the pan). To avoid that problem, CHEF searches the case library for a stir fry plan involving vegetables getting soggy when cooked with meats. It finds a past plan for a beef and broccoli dish that solves this problem by cooking the vegetable and meat separately. The critical similarity between the current situation and the past plan is not the specific ingredients used, but the interaction of one plan (cooking meat) with another (cooking vegetables). CHEF's expertise in this example derives from knowing what similarities to attend to, and knowing what makes a past case especially relevant to a new problem.

As a result, CHEF, through its specific experiences, incrementally improves its ability to predict problems and create plans in response to its own failures. Its attempts at satisfying its own goals, successful or not, result in an improvement in its ability to do so in the future. Thus, the development of expertise in CHEF is dependent upon particular experiences and is bound to the context of previous cases. Experiences can also play a critical role in anticipating the success of possible plans. In the next section, we will examine the role of experiences in recognizing opportunities to satisfy goals.

Opportunistic Planning and Expertise

The CHEF system demonstrated the utility of learning from failures by anticipating and avoiding past failures. But expertise involves another important role for experience: acquiring the ability to foresee and exploit future opportunities to satisfy one's goals. Taking advantage of improved circumstances to solve problems displays the adaptive nature of expertise: Just as a system in a changing world must contend with situations in which plans fail, it must also deal with situations in which it succeeds in unanticipated ways.

In this section, we will look at the problem of anticipating and recognizing planning opportunities. First, we discuss current models of opportunism, and then present a new model, followed by a detailed example. The CHEF planner (Hammond 1989) responds to failure by repairing its current plan and by repairing the knowledge base (its expectations about the results of its actions) which allowed it to create the plan. We now argue that execution-time opportunities can be responded to in a similar way: the planner should exploit the current opportunity and change its expectations to properly anticipate and exploit the future opportunity.

Models of Opportunistic Planning

Our approach builds on two views of opportunism in planning—that of Hayes-Roth and Hayes-Roth (1979), and that of Birnbaum and Collins (1984). Hayes-Roth and Hayes-Roth present the view that a planner should be able to shift to a different planning strategy on the basis of perceived opportunities, even when those opportunities are unanticipated. Their model, called "opportunistic planning," consists of a blackboard architecture (Lesser *et al.* 1975) and planning specialists that capture planning information at many different levels of abstraction. These specialists include domain-level "errand plan" developers (e.g., specialists that know about routes, stores, or conditions for specific plans) as well as more strategic operators (e.g., specialists that look for clusters of goals, and goals with similar conditions). The planner can jump between strategies as different specialists "notice" that their activation conditions are present. For example, in scheduling a set of errands, a specialist that groups errands by location can interrupt another specialist that orders them by priority. In this way, the planner can respond to opportunities to visit nearby locations as they arise.

More recently, Birnbaum and Collins (1984) presented a theory of opportunism where goals are characterized as agents or "demons"—independent processing entities—that have their own inferential power. In this approach, a suspended goal continues to independently examine the ongoing input of objects and events. If circumstances that allow for the satisfaction of the goal arise, this demon mechanism recognizes them and inserts a plan into the current action agenda. For example, consider an agent trying to obtain both food and water in the wild. Suppose the agent decides to suspend the goal to quench his thirst while pursuing the goal to satisfy his hunger. Then, while searching for food, the agent jumps over a stream, and is able to recognize that the stream contains water, and so affords an opportunity to satisfy the suspended goal (to obtain water).

Birnbaum and Collins (1984) argue that the agent is able to recognize the opportunity afforded by the stream because the suspended goal has the demon-based capability to examine the current situation and initiate inferences. This must be the case, they argue, because there can be no way to decide, at the time of suspending a goal, the exact conditions under which it should be reactivated. Birnbaum (1986) argues further that indexing suspended goals by descriptions of the conditions that both signal and allow their satisfaction is an unworkable approach. The problem is that the overall processing system will have to constantly compare the current state of the world to all of the goals waiting in a separate "memory of unsatisfied tasks." However, these arguments may address an overly constrained model of planning at the time of suspension. Instead of instituting a separate list of unsatisfied tasks in memory, a planner could integrate suspended goals into a single memory bank along with all other knowledge. Then, the normal processes that activate information in memory while processing the environment might sometimes also activate associated suspended goals.

Birnbaum and Collins (1984) also raised the problem of having no way to decide, at the time of suspending a goal, the exact conditions under which it should be reactivated. However, one may not need to anticipate the exact conditions for reactivation; instead, a more general category description may be adequate for recognizing types of opportunities. For example, the agent in the example may not have to anticipate coming across "a stream;" instead, simply noting that one should think about the suspended goal whenever one comes across "water" will specify a set of opportunities (a well, a water hole, a stream, a hollow stump) where it would be opportune to reconsider pursuing the thirst goal. Further, opportunism may be imperfect; that is, perhaps it is sufficient to anticipate some features related to opportunities, even at the cost of missing others. This compromise may be worthwhile in order to avoid the heavy computational costs of attempting to draw inferences from each new object to every suspended goal.

In the next section, we propose a model of opportunism that addresses these problems and can provide a better account of how people recognize opportunities during plan execution. We argue that the power of opportunism rests in the ability to reason at the time of goal suspension, before other activities are resumed. At that point, one can devote resources to drawing inferences about the features that indicate opportunities to satisfy the goal. Then, the elaborated inferences about features that signal opportunities can be stored in memory in association with the suspended goal. Later, when these features are perceived in the environment, the pending goal in memory can be automatically reactivated without any further inference required (Patalano, Seifert, and Hammond 1993; Hammond, Converse, Marks, and Seifert 1993). The successful recognition of opportunities will depend on the features chosen to index the suspended goal in memory. Identifying these features involves selecting a specific plan for satisfying the goal from among past plans in memory, abstracting necessary features for successfully executing the plan, and describing the features in a manner easily observed in the environment through normal understanding processes.

The Predictive Encoding Model of Opportunism

Our model of opportunism is based on the ability to anticipate the features in the world that indicate potential opportunities, and to use those features to index the pending goal in memory. At the time of initial encoding, one anticipates features related to circumstances where success would be possible, and then encodes the suspended goal into memory with those features. Then, while pursuing other activities, those same features may happen to be activated; if so, their activation automatically reactivates the suspended goal in memory just when the more favorable conditions have appeared in the world. We call this process predictive encoding, where the features that indicate the relevance of the plan are anticipated and used to index the plan in memory. This anticipation means that no separate or demon-based inference process has to operate; instead, the inference is prepared and stored in memory so that it is simply reactivated by the features perceived in the environment. The steps involved in the predictive encoding model are:

1. Goals that cannot be fit into a current ongoing plan are considered blocked and work on them is suspended.
2. Suspended goals are indexed (associated) in memory with features that are related to potential opportunities.
3. These memory structures are the same ones used to perceive events in the world as the planner executes other plans. No separate "list" of suspended goals is maintained.
4. Later, as elements of memory are activated by features in the world, the goals associated with them are also activated and then integrated into the current processing queue.

In this way, suspended goals are brought to the planner's attention when conditions change so that the goals can be satisfied. In the meantime, the planner can go on to pursue other goals without the need for further reasoning or inferencing regarding the suspended goal.

The success of this process will depend upon the quality of the indices selected for the initial encoding (Seifert, Hammond, Johnson, Converse, MacDougal, and VanderStoep 1994; Johnson and Seifert 1992). We hypothesize that anticipating the circumstances in the world that are related to opportunities to satisfy the goal, and describing them in easily observable features, constitutes the expertise acquired with experience in a domain. Thus, the more one learns about the circumstances related to different types of opportunities, the better one will be able to predict the features and encode them in a way that maximizes successful recognition, and use of opportunities that arise.

For example, consider the goal of opening a locked door. Assume that all reasoning and activity towards accomplishing this goal is suspended while the planner pursues other, perhaps more important goals. Later, if a key is found, its association with locks may facilitate the retrieval of the goal from memory with no special preparation by the planner. But what if a credit card is found instead? Our model suggests that unless you consider the plan of "slipping the lock" at the time of the initial goal suspension, you will be unlikely to notice the credit card as an opportunity to pursue the suspended goal. By preparing such a memory association in advance, it is as if you were placing a "marker" to wait at the "right place" in memory. When these features are observed at a later point, the suspended goal in memory is then activated, and the planner is "reminded" of the goal awaiting pursuit. With careful predictive encoding, no special inference or reasoning process is needed in order to recognize the opportunity; instead, the normal processes occurring while simply recognizing objects in the current environment will automatically activate suspended goals waiting under appropriate indices in memory.

The advantage to this approach is that recognition of later opportunities is made effortless by formulating memory associations so that only normal, default inferencing about the world is required. The inference needed to connect a suspended goal with an opportunity for its satisfaction is prepared in advance, and indexed in memory in terms of readily observable features of the world. So, when initially encoding the "open door" goal, one might reason about plans like "slipping the lock," and determine the particular features of the needed tool; for example, "a thin, flat, yet stiff plane." Then, when one comes across the credit card, or any object fitting the description (such as a plastic knife), one may recognize an opportunity to open the door. Such "smart" indexing provides one solution to the problem of recognizing future opportunities: by anticipating—at the time of encoding—the circumstances that will comprise an opportunity, at least some opportunities will be easy to recognize. By contrast, the demon model of suspended goals predicts no advance preparation of indices; instead, each demon would actively work on-line to connect each perceived object to its triggering conditions. Rather than automatic activation, each demon would need to constantly reason about its potential relationship to each new object as they are perceived.

An Example of Opportunism in Planning

Consider the following example:

> On April 14th, John realized that he needed tax forms in order to prepare his yearly income taxes. Because he was late for work, he had no time to do anything about it. He then goes on to pursue other goals throughout the day. Later, on his way out to lunch, John noticed that he was passing a First of America Bank branch, and recalled that he needed the tax forms. Since he had time, he stopped and picked up the forms and then continued on to lunch.

In this example, John decides to suspend a goal before deciding exactly how to satisfy it; in effect, he says, "I don't have all the resources (e.g., time) to pursue this goal right now." But later, John is able to recognize the conditions that lead to satisfying this goal. How could any processor accomplish this recognition? Our predictive encoding model specifies the following steps.

Suspending Goals

First, the planner suspends the blocked goals by associating them with the elements of memory that describe potential opportunities. In the example, John's goal to possess tax forms is blocked by lack of time. But before suspending the goal in memory, he does some planning about the goal. For novel or infrequent goals, this may involve identifying possible plans to try, resources needed, etc. After thinking for a bit, he recalls that he saw tax forms in the lobby of his bank last year, provided as a convenience for customers. Notice that "bank" may not be the first plan to come to mind for acquiring tax forms; certainly, other plans (like waiting until the forms arrive in the mail, calling the 800 number, or going to the post office) can be generated. But whichever plan he chooses, John is able to prepare for a later opportunity by reasoning out a potential plan and then linking the suspended goal to its preconditions; in the example, this involves "being at a bank."

Then, during the execution of other plans (like going to lunch), the planner naturally observes the effects of its own plans as well as any other events in its world. So later, while John is walking to a restaurant, he sees and recognizes a "First of America" sign as indicating a "bank." This concept in turn activates the prestored association with the suspended tax forms goal. He then recognizes the bank as an opportunity to look for tax forms, and can decide whether to interrupt his lunch to pursue the suspended goal. The key is that no inferencing or reasoning about the possible connection of "bank" to a suspended goal need be performed; instead, the inference is prestored, so that simply recognizing "bank" reactivated the inference made at the time of encoding. Thus, the planner's general recognition of a situation results in immediately activating any goals that have previously been associated with that situation.

What if John had not stopped to try to "predictively encode" the suspended goal? He could have simply realized he did not have time that morning, and then gone on to pursue other goals during the day. But without predictive encoding, would John have recognized the bank as an opportunity? To do so, John would have to examine every new feature of the world and reason about it in order to see if it is somehow related to the tax forms (Is this gas station going to lead to tax forms? What about that flower shop?); in addition, trying to connect everything observed to each of his pending goals would require many comparisons, and could consume much of his cognitive resources. However, by predictively encoding the pending goal, the inferences needed to connect "bank" to the tax forms has already been completed and "canned" into memory; now, all John needed to do was to notice a bank at any time and any place during his day.

Anticipating Opportunities

Predictive encoding requires the planner to anticipate a plan for the suspended goal and the necessary features for executing that plan. But how can a planner best anticipate the critical features of the world that will provide an opportunity for a plan? In general, opportunities to run plans can be derived from the conditions necessary to execute each of the steps of a plan. For example, in order to buy a newspaper, one has to have enough money, be at a store, the store has to have papers in stock, and so forth. So, with enough time, a planner could think through a plan step by step and collect the conditions that have to be true at each point in the plan. This process would require the examination of many conditions, and most may not be particularly helpful in detecting opportunities. For example, some conditions for obtaining tax forms—"have time" and "can carry the forms"—are not very distinctive as features that will allow one to recognize opportunities at the appropriate time. Features that indicate opportunities are more constrained than the conditions necessary for executing a plan; in fact, features indicating opportunities are further constrained by the need for ease of recognition, likelihood of occurrence, distinctiveness, and predictiveness. For example, "have time" is a necessary condition for finding tax forms, but if the suspended goal is associated in memory with time, the planner would be reminded of the goal far too often. While it is a needed condition, it does not isolate situations where acquiring tax forms specifically will be most likely to be successful.

Our solution involves a taxonomy of opportunity types to derive the conditions that will identify opportunities to satisfy goals. The following taxonomy is used to guide the planner's search through a plan to determine appropriate indices to be associated with the suspended goal in memory.

Does the plan require:

- a special resource? If so, index the goal in terms of the resource.
- a special tool? If so, index the goal in terms of the tool.

- a special location? If so, index the goal in terms of the location.
- a special agent or skill? If so, index the goal in terms of the agent or skill.
- a specific time? If so, index the goal in terms of the time constraint.

Using these tests, it is possible to associate the suspended goal in memory (to possess tax forms) with a location (BANK). Other conditions of the plan (such as "possess time") are checked before the plan is actually executed, but they are not used to index the suspended goal in memory. Anticipating the critical features that indicate an opportunity is the core of the predictive encoding process. Recognition of this prepared opportunity may then follow without any special effort, as described in the following section.

Recognizing Opportunities

The predictive encoding model also relies heavily on a model of memory where instances of concepts can be readily identified. For example, John can pick up tax forms at *any* bank, not just a particular one. It is necessary, then, to be able to recognize a wide variety of situations as instances of more general concepts. To accomplish this, we use a version of Martin's (1990) DMAP parser, a general purpose memory recognition system. DMAP uses a marker-passing algorithm in which two types of markers are used to activate and predict concepts (like plans and objects) in an ISA and PART-OF network. Activation markers are passed from primitive features up an abstraction hierarchy. When any part of a concept is activated, prediction markers are spread to its connected concepts, which themselves become active. Of course, this model will only be as good as the network of concepts upon which it operates; in human memory, this conceptual memory network is quite flexible and efficient (Smith and Medin 1981).

For opportunistic planning, we propose adding a new type of link to the memory structures in DMAP. This link associates suspended goals with conditions that represent opportunities to achieve them. Pointing from specific objects to goals, this SUSPEND link is traversed by any activation marker that is placed on the concept. Consequently, the activation of an object also activates any suspended goals associated with it. In the example, the suspended goal to get the tax forms is encoded into memory with the concept representing "you are at a bank." Then, as the world is perceived, a sequence of "parking lot," "building," and "First of America" sign is recognized. Because DMAP is passing activation markers up ISA links, the bank is recognized as a particular First of America branch, an instance of all First of America banks, and of a BANK. While the suspended goal is not directly associated with the concept "First of America," it is associated with BANK, and that causes the activation of the suspended goal to get the tax forms. Accomplishing the activation of the SUSPEND link is thus the same recognition process that occurs during the identification of any perceived object. The power of predictive encoding arises from prestoring the opportunity so that the suspended goal will be activated during normal perceptual processing.

Note that this process will result in sometimes failing to notice an opportunity for satisfying a goal. Some opportunities will be missed when they arise in a form that was not anticipated at the time of the goal suspension. Truly unique opportunities will be missed simply because this technique is aimed at recognition of previously experienced solutions. In addition, opportunities that could have been anticipated, but were not at the time of encoding, will also be missed. For example, if John also passes a library on his way to lunch, he may not recognize it as a potential place to pick up the tax forms. However, if, at the time of the blocked goal, John had anticipated the sources for tax forms, and remembered that libraries and post offices often provide them, he would then be more likely to notice an instance of one of these offices as a potential opportunity to satisfy his goal.

Of course, once a suspended goal is reactivated, it has to be evaluated for integration into the current execution agenda. In the example, the overall plan is changed to take the planner into the bank for a moment to pick up the forms before resuming the trip out to lunch. Successful retrieval and satisfaction of the pending goal serves to confirm and reinforce the predictions made at encoding time. The predictive encoding model as outlined here appears to account for some examples of opportunistic planning. In the next sections, we present some formal tests of the model, first through a computer simulation and then through empirical evidence from psychological experiments.

An Implementation of Opportunism

Our simulations with an implementation of an opportunistic planner were in the TRUCKER program (Hammond, Marks, and Converse 1989). Its domain is a UPS-like pickup and delivery task in which new orders are received during the course of a day. Its task is to schedule the orders and develop the routes for its trucks to follow through town. A dispatcher controls a fleet of trucks that roam a simulated city, picking up and dropping off parcels at designated addresses. Transport orders are "phoned in" by customers at various times during the simulated business day, and the planner must see to it that all deliveries are successfully completed by their deadlines.

TRUCKER's task involves receiving requests from customers, making decisions about which truck to assign a given request, deciding in what order given parcels should be picked up and dropped off, figuring out routes for the trucks to follow, and monitoring the execution of the plans it has constructed. A number of limited resources must be managed, including the trucks themselves, their gas and cargo space, and the planner's own planning time. TRUCKER starts off with very little information about the world that its trucks will be negotiating; all it has is the equivalent of a street map of its simulated world.

TRUCKER's central control structure is a queue-based executor, with planning and monitoring actions sharing space on the queue. The planner must also react to new goals as they arrive. The planner hands pickup and drop-off orders to trucks based on availability in the order they arrive, and integrates the new orders onto each truck's agenda. In the case where the truck is idle, no real integration is necessary, and the standard plan of traveling to the pickup point and then to the drop-off point is used; for example:

```
((GOTO (800 E-61-ST))
(PICKUP PARCEL.42 (850 E-61-ST))
(GOTO (6200 S-COTTAGE))
(DROPOFF PARCEL.42 (6230 S-COTTAGE)))
```

If no idle trucks are available, and if the planner decides the request cannot be merged onto any truck's current agenda, TRUCKER is forced to place the request on a queue of orders waiting for idle trucks. To suspend a goal, TRUCKER marks its representation of the goal's pickup and delivery points with an annotation that there is a goal related to those locations. Because the domain involves only one type of goal, TRUCKER doesn't need to do any more generalization in order to identify good opportunities to satisfy the suspended goals.

Then, as TRUCKER decides when to execute actions based on locations, landmarks, and addresses that it recognizes in the world, it must parse and interpret these objects as they appear in the environment. TRUCKER moves through its world, identifying the objects at its current location and responding with any actions it can now execute. For example, seeing "Woodlawn Street" enables it to execute a plan to turn south on Woodlawn. It also checks for any annotation of a goal with which the object might be associated. If found, the suspended goal is activated and integrated into the current schedule.

At this point the planner is invoked to construct a new route that will satisfy both goals. The fragment of program output in figure 1 illustrates the process of noticing the opportunity, the subsequent reassignment of a request from a different truck, and the construction of a new, combined agenda that includes the opportunity.

After having noticed the opportunity and suggesting the reassignment, TRUCKER uses planning techniques tailored to its domain to plan the new agenda (i.e., specifics about how to combine plans when scheduling trucks). Scheduling the initial pickup is trivial, in that a truck is already at the pickup location. The difficulty lies in scheduling the delivery. TRUCKER does so (though humans may not) by reviewing each location already scheduled and finding the section of the route that will be the least altered by the insertion of the delivery. This can be done even before the exact routes are selected by using the map.

```
8:17:12 AM *** Truck #1 starting new route at 800 block of E 61st Street. ***
*** Truck #1 has noticed opportunity to make pickup for REQUEST.49. ***
*** Request REQUEST.49 is assigned—Truck #1 inserting request in active agenda. ***
*** Noting combination opportunity in memory. ***
8:17:45 AM Planner Action:
                (REASSIGN-BY-NOTICED-OPPORTUNITY REQUEST.49 #1)
— Reassignment of REQUEST.49 means that truck #2 need not continue to destination.
PLANNER assigning REQUEST.49 to truck #1
Starting INTEGRATE-REQUEST REQUEST.49 #1 NOTICED-OPPORTUNITY
Current plan:      ((GOTO (6200 S-COTTAGE) #Structure ROUTE 2)
                    (DROPOFF PARCEL.42 (6230 S-COTTAGE)))
Request-plan to integrate new plan with current plan:
                    ((GOTO (6100 S-COTTAGE))
                    (PICKUP PARCEL.50 (6150 S-COTTAGE))
                    (GOTO (900 E-63-ST))
                    (DROPOFF PARCEL.50 (925 E-63-ST)))
8:17:45 AM *** Truck #1 is stopping in 6100 block of S Cottage Grove. ***
Finishing INTEGRATE-REQUEST REQUEST.49 #1 NOTICED-OPPORTUNITY
Resulting combined plan:
                    ((GOTO (6100 S-COTTAGE))
                    (PICKUP PARCEL.50 (6150 S-COTTAGE))
                    (GOTO (6200 S-COTTAGE))
                    (DROPOFF PARCEL.42 (6230 S-COTTAGE))
                    (GOTO (900 E-63-ST))
                    (DROPOFF PARCEL.50 (925 E-63-ST)))
8:17:48 AM *** Truck #1 making pick-up at 6150 S-COTTAGE. ***
```

Figure 1.

This process for recognizing opportunities has several important features: The full planning mechanisms are invoked only when a suspended goal is activated, suggesting that it is an opportune time to incorporate it into the current agenda. In addition, the same memory for places and landmarks that is used to tell the trucks when to turn and where to stop is annotated with the suspended delivery goals that have not yet been satisfied. Finally, the overhead on this activation is trivial, in that all that TRUCKER has to do is look for a link to a goal on each of the objects it recognizes.

In TRUCKER, the experience of this successful use of opportunity is also saved in memory so that it may be exploited again. In TRUCKER's domain, where a major percentage of the orders are stable over time, certain conjuncts of orders will also tend to be stable and thus predictable. Thus, the plans for those conjuncts can be reused directly, saving both planning time (because the interactions no longer have to be reasoned about) and execution time (because an optimized version of the plan for the conjuncts is being applied). TRUCKER builds a library of routes and conjoined plans for groups of requests that have occurred together. This method of learning conjunctive plans confers three sorts of advantages to TRUCKER:

1. There is the benefit of having learned about profitable combinations of goals in the course of activity that had to be done anyway.
2. There is the benefit of having performed and saved the results of past planning for use in those situations where their utility has been demonstrated.
3. There is the benefit of reduced complexity of the remaining planning and scheduling problem because of the implicit policy decision to reuse combinations that have worked in the past.

However, there are two major limitations with TRUCKER as a testbed for the predictive encoding model. First, although opportunity recognition is crucial, there is really only one type of opportunity that can be exploited in this domain: the presence of a truck in a particular area. Second, although multiple trucks need to be tracked, and multiple orders need to be interleaved, there is not much in the way of interesting variations and interactions among plans in the delivery domain. A domain that is richer in the variety of plans that are executed and in the potential interactions among plans may provide more challenges. If the requirements for plan execution are less similar and more indirect, the need to anticipate opportunities in terms of more general indices will provide a more powerful test of the predictive encoding model.

The success of the TRUCKER program suggests that predictive encoding may be a functional approach to the problem of recognizing opportunities. However, is this approach the same as the one that human planners use? Ideally, we would like to examine the viability of the predictive encoding model as an account of the cognitive mechanisms operative in human experts. In the next section, we present several psychology experiments that address the psychological plausibility of the predictive encoding model. Specifically, the experiments examine the effects of preparing indices at the time of goal suspension on the later recognition of opportunities.

Empirical Evidence for Opportunism

A recent series of studies has addressed the specific predictions of predictive encoding in recognizing opportunities. These studies involved the domain of commonsense planning, where we could presume all of our college student participants are relatively expert. When participants are engaged in a multiple goal planning task, their ability to recognize opportunities to execute suspended goals provides a test of whether predictive encoding is a viable theory of opportunism.

The predictive encoding model previously presented provides several predictions about the recognition of opportunities:

1. that recognition is likely to occur based on features that could be anticipated from planning knowledge about the goal in general;
2. that features associated with opportunities at the time of suspension will be more likely to lead to retrieval of the pending goal;
3. that opportunities not anticipated may be missed;
4. that features requiring inferential links to connect to goals will be less likely to lead to recognition.

The predictive encoding hypothesis thus predicts specifically that opportunities to achieve goals will only be recognized at execution if they have been anticipated at the time of goal suspension.

These predictions from the predictive encoding model of opportunism distinguish it from demon-based models, and if verified, would support it in contrast to other models. Demon-based models, such as Birnbaum and Collins's (1984) "goals as agents" model, predict that all opportunities to achieve goals will be recognized regardless of encoding context. This prediction follows from the fact that the demons are free to make any inferences necessary at the time of exposure to the opportunity. Each goal "agent" can perform as much reasoning and inference as needed to attempt to connect its triggering conditions to the state of the world. Therefore, there is no cost to failing to prepare at encoding; instead, the agent can wait until it is exposed to the opportunity to form this connection. The contrasting predictions from these models are tested by (1) varying the degree to which plans to achieve pending goals are anticipated and (2) comparing whether or not recall cues that have been anticipated in the context of the earlier plans lead to more remindings than do unanticipated cues.

In order to examine these questions about opportunism, we need an experimental paradigm that will allow the creation of goals, exposure to various cues, and then measurement of retrieval of the original goals. This paradigm would allow the direct examination of the cue features leading to appropriate and inappropriate remindings of suspended goals. The method we developed involved a planning scenario including objects likely to be familiar to the college students who serve as our participants. The context of the task scenario for goal presentation was a college dormitory room. Each participant was asked to imagine that he or she must accomplish a number of goals in the dormitory room in a limited period of time. The scenario information provided was as follows:

> Imagine that you are visiting your best friend, Chris, in her dormitory room. After chatting with one another for a while, you both hear a knock at the door. A neighbor peeks her head in and summons Chris to attend a spur-of-the-moment hall meeting. Chris announces she'll be back soon and strolls down the hall to see what's up. In the first few minutes that you are alone in Chris's room, you realize that this is a perfect opportunity for you to do some snooping around. There are all kinds of things that you'd like to know about your friend Chris! And, if you are careful to leave no sign that you've tampered with anything, she'll never find you out.

Within this scenario, a series of goals were presented to the participants. These goal materials were designed with the physical setting in the dorm room in mind; consequently, they were constrained by the objects and activities likely to be possible in this setting. The set of goals presented included the following:

> You notice that Chris left her new college ring on her bureau. You try it on your finger and it gets stuck. Chris will kill you if she finds out that you were so careless with her new piece of jewelry. You need to get the ring off before Chris returns.
>
> You lean over Chris's bed. In the process, you manage to leave scuff marks high up on the white wall next to the bed. Though the marks are faint, they are in a very unusual and conspicuous location. You need to make the wall white again before Chris returns.
>
> When you open Chris's window to get some fresh air, a cold breeze blows her poster off the wall. You are not sure how the poster had previously been attached to the wall, but you do know that you need to reattach it before Chris returns.

Participants were told to read and make a mental note of each goal since they would need to retrieve the goals from memory at a later time. Thus, multiple goals were generated and remembered by the participants, and later retrieved during a cued recall task.

For the measurement of recall, a reminding paradigm was used based on prior studies of retrieving complex episodes from memory (Gentner and Landers 1985; Ratterman and Gentner 1987; Seifert, Abelson, and McKoon 1984). Descriptions of objects that might be found in this dormitory setting were presented to participants as later retrieval cues. The cued-recall task involved presenting a specific description of an object (e.g., "The only thing you find under the sink is a jar of Vaseline. If you could use the Vaseline in a plan to achieve any of your goals, record the plan(s) below"), and asking the participant to write down any prior goals that came to mind. The cue objects presented were assumed to present some salient or basic-level features during the process of comprehending the cue information. This feature information could potentially identify a particular goal previously encoded into memory, and thus retrieve a specific goal and its related information.

The set of cues was developed by identifying two different plans that could satisfy each of the ten goals. Each of these plans was then specified in terms of a specific object that would serve to execute the plan. For example, the goal of removing the stuck ring from the finger could be satisfied by lubricating the finger with Vaseline, or applying cold to the finger to shrink it using ice cubes. One of these two cues—"Vaseline" or "ice cubes"—was presented as a goal-related cue during the cued-recall test, and each was presented to half of the participants. In addition to these ten cues (one for each goal), each participant also saw the same five filler items that were not readily associated with any goal (e.g., comb, tea bags, shoe). As each cue was presented, participants were instructed to consider whether the object listed could be of any help with the goals presented earlier, and if so, to record the goal that "comes to mind." The dependent measure thus consisted of the remindings elicited by these fifteen cues.

The first experiment compared the memory task described above with a nonmemory task. In this No-Memory Control group (each of the groups in the study had over forty participants), all ten goals remained available throughout the session, so they did not need to recall the goals from memory. By comparing this group to a different Memory Control group, we could determine the effects of memory retrieval on the matching task. The reminding data were scored by counting a response as an instance of a reminding whenever it uniquely identified one of the earlier goals. The results showed that the No-Memory group generated 89% of the target goals in response to the object cues, illustrating that participants generally had little difficulty recognizing the intended uses of relevant objects. However, in the Memory Control group, participants generated about 75% of the relevant goals from the target cues. Though this number is less than the No-Memory group, this condition shows that participants were reliably able to recall many relevant goals from memory when cued by target objects.

Two other groups of participants were asked to perform some planning during the goal study phase in order to provide a comparison of goal retrieval when there was advance preparation of plans. In one planning group (the Given Plan condition), each goal was presented with a plan and a specific object for its execution; for example, for the "stuck ring" goal, the plan was "You think that if only you had some Vaseline, you might be able to grease your finger and slide the ring off." Participants in this group simply read the goal and plan information for each goal. In another planning group, the Guided Plan condition, participants were given a specific object with each goal, and were asked to generate the plan; for example, "You think that if only you had some Vaseline, you might be able to...". In this condition, the participants may work more actively to link the object with the goal to be achieved. In fact, all plans generated by participants during this phase were consistent with those intended by the experimenter (i.e., all plans matched those given in the Given Plan condition).

The mean percentage of target remindings generated in response to rehearsed versus unrehearsed cues was computed for each of the conditions. As expected, for both the Given Plan and Guided Plan conditions, a greater percentage of target goals were generated in response to rehearsed (89%) as compared with unrehearsed (71%) cues. A participant who, for example, associated Vaseline with the goal of removing a stuck ring and tape with the goal of re-hanging the fallen poster in the study phase of the experiment was more likely to be reminded of the stuck ring upon seeing "Vaseline" than to be reminded of the poster upon seeing "masking tape." This finding supports the hypothesis that associating a plan with a goal at the time of encoding facilitates opportunistic remindings.

Rehearsed cues not only led to a greater number of remindings than unrehearsed cues, they also led to a decrease in the mean number of unintended remindings generated for each cue. A greater mean number of unintended remindings was generated in response to each unrehearsed cue (32%) versus each rehearsed cue (20%) across Given and Guided Plan conditions, with an even higher rate of unintended remindings to the filler cues (44%). These differences suggest a task demand characteristic: that participants attempted to write a response to each cue, even if the response was not an appropriate one.

The study suggests that participants were more often reminded of goals when faced with anticipated cues because they had previously encoded each goal in terms of the cue. When processing the cue, the goal came to mind because it had previously been associated with the cue in memory. This predictive encoding effect is reminiscent of Tulving and Thomson's (1973) encoding specificity effect, where retrieval was shown to depend upon encoding. However, in our study, it was not simply *any* commonality between encoding and retrieval contexts that led to increased recall. Instead, the rehearsal of specific plan-related information, specified by the predictive encoding model, led to remindings from later cues. This finding contradicts a prediction from a demon-based model; if the inferences that connect an opportunity to a suspended goal are made at the time of exposure to the object, then differences in preparation should have little effect on recognition. But the results show that the specific associations formed at the time of suspension accounted for differences in recognizing opportunities.

It is possible that even less specific associations may serve to retrieve the pending goal at a later time; for example, recognition of opportunities that fit the same plan but are not instantiated in a particular specific object. Optimally, planning preparation could be more general and still produce appropriate remindings. In a second experiment, rather than rehearsing a specific plan and object, participants were either given or were asked to generate a general type of plan that would satisfy each goal. For example, in the Given Plan condition, the plan included only a type of plan: "You think that you might be able to lubricate your finger and slide the ring off." In Experiment 2, a "Guided" planning condition was not possible, since specific objects were to be avoided in the descriptions. Instead, a Generate Plan condition was added, wherein participants were asked to devise their own type of plan for each goal, without specifying a particular object that might be used to accomplish it. The instructions for plan generation made explicit the need to be general: "Keep in mind that we are not interested in what specific objects you would use to resolve each situation. Rather, we are interested in the kinds of approaches you would take. So, for example, you need not write plans like 'Use a sledgehammer to break down the door' since you could use any one of a number of objects to do so." In the Generate Plan condition, participants generated about two plans for each goal; for each participant, one of the plans generated was related to later reminding cues for around six of the ten goals.

The number of target remindings generated in response to rehearsed versus unrehearsed cues showed that, as expected, for the two conditions where advance planning occurred (the Given Plan and Generate Plan conditions), a greater percentage of target goals were recalled in response to rehearsed (82%) as compared with unrehearsed (65%) cues. Other measures, such as a free recall baseline at the end of the experiment, refute the possibility that participants were reminded of a greater percentage of goals associated with rehearsed cues simply because these goals were more available in memory. In fact, there were no differences in the groups or rehearsal conditions in overall free recall of the goals in memory.

A third experiment repeated the encoding of goals in association with potential plans to achieve them. In this study, all of the participants performed the Guided Planning task for all ten goals; for example, in studying the goal to remove the ring, they were given the object "Vaseline" and asked to generate a plan using that object to achieve the goal. But instead of seeing the identical object or a novel object as a recall cue, participants sometimes saw an object related to the specific plan they had studied. For example, in the cued reminding portion, participants who studied "Vaseline" saw either an unrehearsed cue ("ice cubes"), the rehearsed cue ("Vaseline"), or a different object that also fit the rehearsed plan ("butter"). (Note that all three cues are objects that can be used with some plan to satisfy the goal.) This manipulation will tell us whether the advantages of rehearsal are specific to the object studied, or whether participants spontaneously generated the more general plan category needed for the goal ("a kind of lubricant").

The results showed that all three types of cues resulted in more appropriate remindings than did unrelated cues. As expected, the cue actually rehearsed resulted in the most remindings (98%), while the unrehearsed related cue resulted in fewer (63%). However, the alternative version of the rehearsed cue ("butter") resulted in more remindings (76%) than the unrehearsed related cue, showing that rehearsal of a specific cue facilitated remindings from a more general version of that rehearsed cue. This finding suggests that predictive encoding does result in advantages to the planner in recognizing later opportunities: Not only is the rehearsed cue reliably tied to the suspended goal, but other unseen objects related to the rehearsed plan also result in more remindings than would occur without any rehearsal. So, attempting to encode the features related to potential opportunities benefits recognition not only of the specific case, but also more generally of items related to the rehearsed plan.

These results illustrate the advantages of preparing for later opportunities: when a participant had anticipated the type of plan for the goal, objects related to that plan were sufficient to enhance retrieval of the goal compared to unrehearsed objects. Associating a plan object with a goal *at the time of encoding* facilitated opportunistic remindings, but anticipating a type of plan more generally was also sufficient. Further, encoding specific rehearsed plans and objects resulted in remindings based on later objects that were only abstractly related to the plans studied. This finding raises questions about what level of generality is best for preparing features, because optimally, one would not want to be too general or too specific in preparing indices to goals in memory (and consequently be reminded too frequently or too infrequently to be of value). Further studies are pursuing this optimality issue through both experimental and modeling efforts.

In summary, these empirical studies support the predictive encoding hypothesis, in which planning needs are anticipated at the time of goal suspension, the goal is indexed with readily-observed features, and new experiences then retrieve the suspended goal from memory.

Conclusions

Opportunistic planning may represent a satisfying solution to the problem of how much to preplan efforts to achieve goals (as in complete planners) vs. waiting for the world to provide solutions to unsatisfied goals (as in situated approaches). The predictive encoding approach involves some preplanning to anticipate likely opportunities that may arise; however, it also involves reacting to the environment by allowing happenstance to activate pending goals. This predictive encoding solution is "intelligent" in that it takes advantage of both our ability to plan and our inability to predict the changes that might occur in the world we experience. Its consequence is that we may still miss opportunities that were potentially observable, while still maximizing the detection of opportunities that we know, from experience, are likely to occur.

Predictive encoding is important because it suggests that how people encode goal-related information is extremely important to their being able to recognize opportunities to achieve their goals. And recognizing opportunities is of great value when planning in a dynamic world. Given the continuous changes in circumstances as one moves through the environment, and the multiple-goal nature of most pursuits, it is not often possible to be aware of all of the features of the environment, let alone to stop and reason about the relevance of each to all pending goals. Predictive encoding allows a planner to retrieve goals from memory "on demand," while avoiding the need for heavy inferential processing of each new stimulus. Its functional utility for machine planners and its consistency with psychological evidence suggests predictive encoding is a viable account of a mechanism for opportunistic planning.

This work also suggests a mechanism for expertise to affect performance on other types of cognitive tasks. The ability to predict the circumstances leading to the solution of goals appears central to recognizing opportunities, but it may also affect success in generating potential diagnoses, avoiding errors, and in learning from experience. Perhaps this ability to project potential solutions helps the expert gain from experiences though comparing a strong "hypothesis" based on existing knowledge to the actual outcome of experience in the world. In this way, knowledge refinement and further learning may be driven by the active process of predicting potential paths to goal success. This suggests that specific experiences within a new domain may provide a critical training ground for the development of expertise, where experts can learn to anticipate and take advantage of opportunities to achieve their goals. And it is through ongoing experiences of attempting to satisfy goals in a dynamic world that opportunities for developing expertise can be encountered.

Acknowledgments

This work was supported by the Office of Naval Research under contracts N0014-88-K-0295 and N00014-91-J-1185 to the University of Chicago and contract N00014-91-J-1128 to the University of Michigan. Portions of this chapter have previously appeared in the *Journal of Machine Learning*, and some of the empirical results were first presented at the Cognitive Science Society meeting in 1994.

References

Bareiss, R. (1989). *Exemplar-Based Knowledge Acquisition.* San Diego: Academic Press.

Birnbaum, L.A. (1986). Integrated processing in planning. Ph.D. dissertation, Yale University, New Haven, CT.

Birnbaum, L.A & Collins, G.C. (1984). Opportunistic planning and Freudian slips. *Proceedings of the Sixth Cognitive Science Society,* Boulder, CO.

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

Carbonell, J. (1981). A computational model of analogical problem solving. *Proceedings of the Seventh International Joint Conference on Artificial Intelligence.*

Gentner, D. & Landers, R. (1985). Analogical reminding: A good match is hard to find. *Proceedings of the International Conference on Cybernetics and Society* (pp. 607-613), Tucson, AZ. New York: IEEE.

Hammond, K.J. (1989). *Case-Based Planning: Viewing Planning as a Memory Task.* San Diego: Academic Press.

Hammond, K.J., Marks, M. & Converse, T.M. (1989). Planning in an open world: A pluralistic approach. *Proceedings of the Eleventh Annual Cognitive Science Society,* Ann Arbor, Michigan.

Hammond, K.J., Converse, T.M., Marks, M. & Seifert, C.M. (1993). Opportunism and learning. *Machine Learning, 10* (3), 279-310.

Hayes-Roth, B. & Hayes-Roth, F. (1979). A cognitive model of planning. *Cognitive Science, 3,* 275-310.

Johnson, H.M. & Seifert, C.M. (1992). The role of predictive features in retrieving analogical cases. *Journal of Memory and Language, 31,* 648-667.

Kolodner, J., Simpson, R. & Sycara, K. (1985). A process model of case-based reasoning in problem-solving. *Proceedings of the Ninth International Joint Conference on Artificial Intelligence,* Los Angeles, CA.

Lesser, V.R., Fennell, R.D., Erman, L.D. & Reddy, D.R. (1975). Organization of the Hearsay-II speech understanding system. *IEEE Transactions on Acoustics, Speech and Signal Processing, ASSP-23,* 11-23.

Martin, C.E. (1990). *Direct memory access parsing.* Ph.D. dissertation, Yale University, New Haven, CT.

Patalano, A.L., Seifert, C.M. & Hammond, K.J. (1993). Predictive encoding: Planning for opportunities. *Proceedings of the Fifteenth Annual Cognitive Science Society Conference* (pp. 800-805), Boulder, CO.

Ratterman, M.J. & Gentner, D. (1987). Analogy and similarity: Determinants of accessibility and inferential soundness. *Proceedings of the Ninth Annual Meeting of the Cognitive Science Society* (pp. 23-24). Seattle, WA.

Schank, R. (1982). *Dynamic Memory: A Theory of Learning in Computers and People.* Cambridge: Cambridge University Press.

Seifert, C.M., Abelson, R.P. & McKoon, G. (1984). Being reminded of thematically similar episodes. *Proceedings of the Sixth Conference of the Cognitive Science Society* (pp. 310-314), Boulder, Colorado.

Seifert, C.M., Hammond, K.J., Johnson, H.M., Converse, T.M., MacDougal, T. & VanderStoep, S.W. (1994). Case-based learning: Predictive features in indexing. *Machine Learning, 16,* 37-56.

Smith, E.E. & Medin, D.L. (1981). *Categories and Concepts.* Cambridge, MA: Harvard Press.

Sussman, G. (1975). *A Computer Model of Skill Acquisition,* Vol. 1 of Artificial Intelligence Series. New York: American Elsevier.

Tulving, E. & Thomson, D. (1973). Encoding specificity and retrieval processes in episodic memory. *Psychological Review, 80,* 352-373.