

# Predictive Encoding: Planning for Opportunities\*

**Andrea L. Patalano and Colleen M. Seifert**

Department of Psychology  
University of Michigan  
330 Packard Road  
Ann Arbor, MI 48104  
(313) 763-0008  
patalano@psycho.psych.lsa.umich.edu  
seifert@um.cc.umich.edu

**Kristian J. Hammond**

Department of Computer Science  
The University of Chicago  
1100 East 58th Street  
Chicago, IL 60637  
(312) 702-1571  
hammond@tartarus.uchicago.edu

## Abstract

Suspended goals are those that are postponed by an agent because they do not fit into the agent's current, ongoing agenda of plans. Recognizing later opportunities to achieve suspended goals is an important cognitive ability because it means that one can defer work on a goal until one is in a better position to achieve the goal. This paper focuses on when and how such opportunities are recognized in everyday planning situations. According to our account of the phenomenon, suspended goals are associated at the time of encoding with features of the environment in which goal achievement would likely be possible. This process is referred to as predictive encoding. Later, when these features are perceived in the environment through normal inferential processes, the agent is reminded of suspended goals through features previously associated with them, and recognizes the opportunity to achieve the goals. This approach is compared with other recent theories of opportunistic planning, and empirical work is presented which supports predictive encoding as an explanation for opportunistic planning behavior.

---

\* This work was supported by the Office of Naval Research under contract N0014-85-K-010 to the University of Chicago and contract N0014-91-J-1128 to the University of Michigan.

## Introduction

Consider the following example:

On making breakfast for himself in the morning, Bob realized that he was out of orange juice. Because he was late for work, he had no time to do anything about it. Later, on his way home from work, Bob noticed that he was passing a "Stop and Shop," and recalled that he needed orange juice. Having time, he stopped and picked up a quart and then continued home.

This example illustrates a rather mundane, everyday experience. One is often reminded of an unrealized goal at the moment that he or she is in an opportune position to achieve it, just as Bob was reminded to buy orange juice upon driving past a grocery store. However, from the viewpoint of explaining the mechanisms involved in human cognition, Bob's behavior suggests the intriguing question of *how* such a situation is recognized as an opportunity to achieve a previously-formulated goal.

The term "suspended goal" (Schank & Abelson, 1977) refers to a goal that cannot be fit into current, ongoing planning activity, and is postponed rather than abandoned. Bob's buying orange juice was a suspended goal because it did not fit into the current plan of getting to work on time, but was put off to be pursued at some later point. Recognizing

opportunities to achieve suspended goals is an important cognitive ability because it allows one to defer work towards a goal until one is in better circumstances to achieve the goal. In other words, suspending goals allows one to plan *opportunistically*, creating the most efficient plans and execution schedule for realizing a set of goals in a dynamic environment.

How is it that important goals might come to mind at the appropriate times? A number of models have been suggested to account for opportunistic behavior. Hayes-Roth and Hayes-Roth (1979) offered one of the first of such models which was based on protocols of subjects engaged in errand planning. Hayes-Roth et al.'s computer model, based on a blackboard architecture (Lesser et al., 1975), incorporated their observations that planning is typified by one's ability to shift between planning strategies as unanticipated opportunities are perceived. For example, in scheduling a set of errands, they noticed that planners were able to jump back and forth between ordering errands in terms of location, and ordering them in terms of goal priority. These strategies were represented in their model as "specialists" that gathered information at various levels of abstraction, waiting for their activation conditions to be present. Unfortunately, Hayes-Roth et al.'s view of opportunism includes no memory component. Their subjects neither needed to remember their goals nor to respond to unanticipated execution-time opportunities to achieve them. All of the goals were made available to the subjects throughout the experiment. As such, their model was one of how plans are created to achieve active goals rather than one of how goals that have been stored in memory are brought to mind and achieved in response to execution-time opportunities.

One model of how goals might be suspended in memory for later execution was developed by Marks, Hammond, & Converse (1989) and implemented in TRUCKER, a computer planner. TRUCKER's domain is a UPS-like pickup and delivery task in which new orders are received during the course of the day's execution. Its task is to schedule orders and to develop routes for its trucks to follow through town. When TRUCKER receives goals that it cannot immediately accomplish, the goal becomes "suspended" and is added to the list of pending goals. For example, if it receives a request to pick up a package at the Sears Tower, but no trucks are near the Sears Tower, the goal will be added to the list of goals that need to be accomplished. How does TRUCKER later recognize opportunities to achieve the suspended goal? At the time of encoding the goal of going to the Sears Tower, TRUCKER connects the goal with its internal representation of the Sears Tower. As a result, when a truck passes the Tower, its perceptual processing of the Tower automatically

activates the goal of stopping at the Tower and retrieving the package. In this way, TRUCKER provides a simple model of opportunistic memory. Unfortunately, this model is not complete because it deals only with goals whose execution conditions are well specified. The goal of stopping at the Sears Tower can only be executed when passing the Sears Tower. People, however, generally have goals that can be achieved in multiple ways. The goal of determining what time it is, for example, can be achieved by looking at a clock, turning on the television or radio, or calling an information phone number. The model instantiated in TRUCKER does not explain how the more complex goals faced by people would have to be encoded in memory so that opportunities to achieve them might later be recognized.

A model by Birnbaum and Collins (1984) does suggest a mechanism for recognizing multiple execution-time opportunities to achieve previously-suspended goals. Their approach views goals as independent processing entities that have their own inferential power, similar to "demons." Thus, when a goal is suspended in memory, it acts as a processor that continues to examine the ongoing flow of objects and events that occur. If circumstances relevant to its satisfaction arise, the goal demon recognizes its potential use of these circumstances and becomes active in modifying the current plan to take advantage of the opportunity. As evidence for their approach, Birnbaum and Collins offer a simple yet compelling description of a person trying to obtain both food and water in the wild. In the example, the person decides to suspend the goal of satisfying thirst while trying to achieve the goal of satisfying hunger. While searching for food, however, the person jumps over a stream, and is able to recognize that the stream affords an opportunity to satisfy one suspended goal. If anticipation were one's only means of recognizing opportunities, the opportunity to obtain water would have been missed by the planner in this story. Thus, Birnbaum and Collins argue that the suspended goal must have the ability to examine the current situation and initiate inference to test its own relevance. By positing independent agents responsible for each goal, they provide a mechanism for goal-directed inference.

This approach to opportunistic behavior is an unlikely explanation of human cognitive processes simply because of its computational demands. Given the number of stimuli in the environment at any point in time and the number of goals one might intend to resolve, it is not clear how goals could actively decide on the relevance of each stimulus in relation to each goal. The process might be feasible if all goals were as simple as "finding water in the wilderness," and all features of the environment were as obvious as "a stream," where little inferencing is necessary to

recognize that a stream is an example of water. However, consider again the example of determining the current time. One could call an information hotline to get the time, turn on the television or radio, or start the car. None of these objects, however, are obviously related to finding out the time. A demon would have to do a great deal of inferencing to determine that a car might be used in such a plan.

An additional difficulty with the Birnbaum and Collins model is that it assumes that people recognize all opportunities to achieve suspended goals. While no previous studies have provided contradictory evidence, intuition and personal experience suggest that we often miss opportunities to achieve our goals. It seems that any model that claims to account for human behavior should account for missed opportunities as well as recognized ones.

We suggest that opportunism does not result from goal demons constantly monitoring the world; instead, we argue that the indexing of suspended goals in memory provides a better account of when opportunities will be recognized during plan execution (Hammond et al., in press). According to this model, only normal inferential processes are used to monitor the world. In order for an opportunity to achieve a goal to be recognized, the goal must be stored in memory in terms of features of the environment that are normally being monitored. Recognizing a stream as an opportunity to achieve the goal of finding water is easily explained by this model, since both the goal and the stimulus are described in terms of the same features. If one encodes a goal as "find water" and later identifies a stimulus in the environment as a "stream of water," a direct connection is made between the object and the goal, and the goal is brought to mind. What if the goal were instead "determining the time of day?" According to Birnbaum and Collins, the presence of a clock and a telephone should be equally likely to remind one of the goal. Our theory, however, predicts that while a clock would likely ordinarily serve as a reminder, a telephone would only serve as a reminder if one had previously encoded the goal in terms of the plan "use a telephone to call an information number." Otherwise, the goal "get time" and the telephone, which might be identified as "instrument for making phone calls," would not be linked in memory and no reminding would take place.

Predictive encoding stresses the importance of encoding blocked goals in memory in such a way that they will be recalled by conditions favorable for their solution. If a goal is to be recalled at an opportune time, the features of situations in which goal achievement would be possible must be anticipated, and the suspended goal must be indexed in memory in terms of those features. Indexing means "preparing" connections to a blocked goal so that it is

likely to be retrieved when the features of a new situation match the features associated with the goal in memory. In other words, suspended goals recognize opportunity by waiting for them in the "right place" in memory; if those features associated with the goal are later activated, the goal will also be activated. This process of using memory to prepare for recognizing later opportunities is referred to as predictive encoding because the features that indicate the relevance of a plan are anticipated, and used to index the plan in memory.

The following study investigates recognizing opportunities in human planning. It contrasts the predictive encoding hypothesis that opportunities to achieve goals will only be recognized at execution if anticipated during planning, with Birnbaum and Collins's "goals as agents" prediction that all opportunities to achieve goals are equally like to be recognized regardless of encoding context. These contrasting predictions are tested by varying the degree to which plans to achieve pending goals are anticipated by subjects, and comparing whether or not recall cues that have been anticipated in the context of the earlier plans lead to more reminders than unanticipated cues. Our expectation is that anticipated cues will remind subjects of their goals more often than will unanticipated cues.

The context for goal presentation in the study is a college dormitory room. The subject is asked to imagine that he or she must accomplish a number of goals in the dormitory room in a limited period of time. The goals that must be accomplished, which include re-hanging a fallen poster, retrieving an item from a high shelf, and removing a stuck ring from one's finger, are intended to be plausible for such a setting, and achievable with the aid of objects likely to be found in such a setting. Additionally, all objects presented to subjects as later retrieval cues are objects that might be found in this setting.

## Method

### Subjects

Subjects were 128 undergraduates (65 men and 63 women) at the University of Michigan who participated in partial fulfillment of introductory psychology course requirements. Subjects were assigned to one of three conditions: No-Plan in which subjects read a set of goals, Given-Plan in which subjects also read a suggested plan for achieving a goal, or Generate-Plan in which subjects were allowed to anticipate plans for achieving goals. Subjects were run in groups of approximately twenty in 50 minute sessions.

## Materials and procedure

All subjects were given a test booklet to read and complete. The booklet began with an instruction sheet in which subjects were told that they would be asked to imagine themselves in the following scenario:

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.

Subjects were then told that they would be presented with a set of goals related to the scenario, with one goal being presented on each of the subsequent pages of the workbook. The instructions made clear that subjects should read and make a mental note of each goal since they would need to retrieve the goals from memory at a later time. Each goal was a description of something that, hypothetically-speaking, needed to be accomplished by the subject in the dorm room before the occupant of the room returned. An example goal follows:

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.

All goals were tasks that could be accomplished using objects found in a typical dormitory room. For example, the above goal might be accomplished by rubbing soap around the ring and attempting to slide the ring off. This segment of the booklet will be referred to as the goal learning phase of the experiment.

Following goal learning, subjects were told to imagine that as they continued inspecting the dorm room, they encountered numerous objects, some of which might be useful to them in achieving their goals. They were told that one encountered object would be described on each of the subsequent pages

of the workbook. Subjects were instructed to read about each object and, if they believed the object could be used to achieve an earlier goal, to record in the space provided their plan for how this could be accomplished. One object description was: "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."

Fifteen objects were presented to each subject, 10 of which were intended to suggest a plan to achieve an earlier goal, and 5 of which were intended as fillers. Two different goal-related object sets were used, with filler objects the same for both sets. All materials were pilot tested on an independent group of subjects in a non-memory experiment to ensure that presented objects were seen as components of reasonable plans to achieve the given goals. This segment of the booklet will be referred to as the test phase of the experiment.

The 3 conditions varied in terms of the type of processing performed on the goals. The No-Plan condition was intended as a baseline condition and no processing of the goals was performed beyond what has already been described. In the Given-Plan condition, test booklets followed the same format with one addition: each goal description during goal learning concluded with a sentence suggesting a single plan that might be used to achieve the stated goal. For example, appended to the goal of removing a ring from one's finger might be the statement: "You think that if only you had some Vaseline, you might be able to grease your finger and slide the ring off." Half of the plans seen by each subject suggested an object that later appeared in the test phase for that subject. The remaining plans suggested objects that were not later presented. Thus, objects suggested during goal learning that did not correspond to objects presented during testing, were drawn from the object set that was not used in the test phase.

Finally, in the Generate-Plan condition, unlike the Given-Plan condition, no plans were suggested during goal learning. Rather, in addition to reading about each goal, subjects were instructed to generate plans on their own and to record them in the space provided on each page. Subjects were told to adhere to the following guidelines: to record plans in the order they come to mind, to be specific in their plans, and to be sure to record the objects that would be needed to carry out each plan. Additionally, they were told to generate only realistic plans, i.e. those that could be executed using objects commonly found in a college dormitory room.

## Results

In order to determine the effects of anticipating the plan opportunities, all target cues were first

categorized as anticipated or unanticipated. In the No-Plan condition, all cues were unanticipated in the sense that they were not presented during goal learning, while in the Given-Plan condition, exactly half of the target cues presented to each subject were anticipated (5 were anticipated and 5 were unanticipated). In the Generate-Plan condition, whether or not a cue had been anticipated was determined by comparing each cue to the goal learning plans generated by the subject for the intended associated goal. If the cue was stated in a plan to achieve the goal, it was coded as anticipated. Otherwise, it was coded as unanticipated. Coding was conservative in the sense that if a subject-generated plan could have been achieved using a cued object, but the plan did not specifically mention that object, the cue was coded as being unanticipated. Across all subjects in the Generate-Plan condition, a mean of 3.22 cues were anticipated and 6.78 cues were unanticipated.

Additionally, all test cue responses were coded in terms of whether or not the cue led to the appropriate goal reminding. A reminding was scored whenever the written response could be uniquely identified as an intended earlier goal. Of primary interest was the relative number of goal reminders generated in response to anticipated versus unanticipated recall cues in each of the anticipation conditions (Given-Plan and Generate-Plan conditions). In the Given-Plan condition, a mean proportion of .77 ( $SD = .22$ ) of unanticipated cues led to appropriate reminders, as did .81 ( $SD = .21$ ) of the anticipated cues. In the Generate-Plan condition, proportions of .74 ( $SD = .23$ ) and .96 ( $SD = .13$ ) unanticipated and anticipated cues respectively lead to appropriate reminders. A 2x2 ANOVA revealed a main effect of anticipation ( $F(1, 83) = 27.21, p < .001$ ) across the two conditions and an interaction effect between condition and anticipation ( $F(1, 83) = 12.87, p = .001$ ). No main effect of condition was found ( $F(1, 83) = 2.61, p = .110$ ).

The No-Plan condition was run as a baseline condition to establish the mean number of reminders given no anticipation. One might imagine that an increase in attention to goals during goal learning (due to additional reading or to plan generation) would lead to better memory for all of the goals and thus to an increase in reminders even in response to unanticipated cues. Comparing the reminding means for unanticipated cues in the No-Plan (.74,  $SD = .19$ ), Given-Plan (.77), and Generate-Plan (.74) conditions, it becomes clear that this was not the case ( $F(2, 125) = .25, p = .783$ ). That is, no differences were found across the 3 conditions in terms of the memorability of goals in response to unanticipated cues. To the extent that the manipulations influenced reminders, they appeared to influence reminders only for anticipated cues, confirming that it was anticipation

of the specific cues and not just anticipation in general that lead to increased reminders.

## Discussion

The results illustrate that, as predicted, reminding is more likely when a cue had been anticipated than when it had not. This effect held across both anticipation conditions. The study supports the predictive encoding hypothesis and suggests that the reason subjects were more often reminded of goals when faced with anticipated cues was because they had previously encoded each goal in terms of the associated cue. When processing the cue, the goal came to mind because it had previously been associated with the cue in memory. This effect might be similar to that described by Tulving and his colleagues (Tulving & Thomson, 1973) as encoding specificity whereby how and when a memory is retrieved depends on how it was stored. While Tulving's work on encoding specificity was primarily performed with word lists of paired-associates, the concept of encoding specificity might also be relevant for plans and goals.

Surprisingly, the results of this study also showed that the difference in proportion of reminders in response to anticipated versus unanticipated cues was greater when subjects generated their own plans than when plans were generated for them. One explanation for the interaction is that generating a plan oneself leads to a stronger memory association between the plan and the goal than would be created by one's simply reading about a plan. Given that subjects were instructed only of the importance of attending to each goal (but not of attending to associated plans) in the Given-Plan condition, it is possible that weaker associations were formed between the plan and the goal in this condition. Unfortunately, a second difference between the two conditions was that in the Given-Plan condition, the experiment determined which plans would be anticipated, while in the Generate-Plan condition, subjects determined which plans they would consider in advance. While this limits the conclusions that can be drawn, it motivates future investigation to distinguish the extent to which subject-generation of plans versus degree of encoding of plans (whether they be subject-generated or not) is primarily responsible for facilitating the recognition of opportunities.

An additional consideration raised by the study is the extent to which retrieval cues suggest plans to achieve particular goals. In this study, while all object cues given to subjects were components of reasonable plans to achieve goals, some clearly suggested prior goals more than did others. For example, consider the objects "thumbtack" and

"gumball" as components of plans to achieve the goal "re-hang fallen poster." In processing "thumbtack," the function "can be used to hang papers to a board" would likely come to mind quickly. For "gumball," on the other hand, the function "can be chewed" might be most salient. Given a thumbtack, drawing the inferences needed to get from "can be used to hang paper to a board" to "use now to hang poster to the wall" would be rather straightforward for most people. However, given a gumball, getting from "can be chewed" to "use now to hang poster to wall" would require that a great many more inferences be made (i.e., that gum can be chewed, that chewed gum is sticky, and that sticky material can be used to re-hang dormitory room posters). It is plausible that the degree to which anticipation influences reminding depends on the nature of the inference that must be made in order to recognize the relationship between a cue and a goal. In other words, when a cue directly suggests a goal, rehearsal might have less of an effect on reminding. However, when the path between a plan and a goal is less direct, rehearsal might play a much greater role.

Predictive encoding as a theory is important because it suggests that how people encode goal-related information is extremely important to their being able to recognize opportunities to achieve goals. Our current research in both artificial intelligence and psychology is directed towards resolving some of the theoretical controversies surrounding the mechanisms underlying opportunistic behavior, and to developing a deeper understanding of how and when opportunities are recognized and taken advantage of in everyday planning. Recognizing opportunities is an important component of achieving goals in a dynamic world because we are often unaware of all of the features of the environment, and thus cannot rely on implementing even carefully-constructed plans as a means of achieving all of our goals. Rather, we must not only be able to take advantage of opportunities when they arise, but we must also use this new information to help us predict and plan for future opportunities. Predictive encoding suggests one mechanism by which opportunistic planning may occur.

## References

- Birnbaum, L. A., & Collins, G. C. (1984). Opportunistic Planning and Freudian Slips. Proceedings of the Sixth Cognitive Science Society, Boulder, CO.
- Hayes-Roth, B., & Hayes-Roth, F. (1979). A cognitive model of planning. *Cognitive Science*, 3, 275-310.
- Hammond, K., Converse, T., Marks, M., & Seifert, C. (in press). *Journal of Machine Learning*.
- 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.
- Marks, M., Hammond, K., & Converse, T. (1989). Planning in an Open World: A Pluralistic Approach. Proceedings of the Eleventh Cognitive Science Society, Ann Arbor, MI.
- Schank, R. & Abelson, R. (1977). *Scripts, Plans, Goals and Understanding*. Hillsdale, NJ: Lawrence Erlbaum.
- Tulving, E., & Thomson, D. (1973). Encoding specificity and retrieval processes in episodic memory. *Psychological Review*, 80, 352-373.