A Functional Perspective on Reminding*

Tom McDougal
and
Kristian Hammond
Department of Computer Science
University of Chicago
1100 E. 58th Street
Chicago, IL 60637

Colleen Seifert
Department of Psychology
University of Michigan
330 Packard Road
Ann Arbor, MI 48104

Abstract

This paper explores the relationship between human activity and remindings. We argue that the type of activity in which a person is engaged influences the kinds of features that trigger remindings, and, conversely, that remindings can change situated behavior over time. We argue further that the types of indices used in memory are not uniform, but depend upon the nature of the task that a person is engaged in as well as his history of interaction with the world.

Introduction

A rift currently divides the field of analogical reasoning, with researchers in psychology on one side and those in Artificial Intelligence on the other. This rift comes from a basic difference in research philosophies within the two camps. Researchers in AI generally try to solve problems; they determine the functionality they need, develop theories of behavior to provide that functionality, and then correlate those theories with human behavior. Researchers in psychology, work in the other direction. They try to progress from observed human behavior to a functional theory.

Specifically, a debate has long brewed with regard to people's ability to be reminded on the basis of abstract features. The AI functionalists have constantly promoted the need to find a "best match" between a new problem description and an existing example in memory (Kolodner & Simpson 1989, Hammond 1989, Bareiss 1989). They have concentrated on examples that push the limits of their ability to *represent* relevant features, let alone organize memory around them (Birnbaum and Collins 1989, Schank 1982).

Conversely, the empiricists in psychology claim that the experimental data shows that retrieval occurs only on the basis of surface features (Gentner & Landers 1985, Ross 1987, Gick & Holyoak 1983, Holyoak & Thagard 1985).

This either/or debate hides the diversity of remindings. Memory serves many masters, and thus we should *expect* it to be sensitive to the activity in which the agent is engaged. Sometimes remindings occur on the basis of surface features, sometimes on the basis of deep structure. Evidence exists for each, and each has its unique functionality. Likewise, learning from remindings is not always the one-shot deal often postulated by those in case-based reasoning, but can arise from a dynamic interaction between repeated remindings and activity in the world.

In this paper, we will present several examples of remindings in different contexts, examining each in detail to see what functional role the remindings play, and what they tell us about the structure and dynamic of memory. Our argument will be that, while memory is driven by a single process, the different functions it serves cause it to exhibit different behaviors.

* This work was supported in part by the Defense Advanced Research Projects Agency, monitored by the Air Force Office of Scientific Research under contract F49620-88-C-0058, and the Office of Naval Research under contract N0014-85-K-010.

Two Examples of Remindings

We begin with two examples which address the abstract feature/surface feature dispute, with implications for planning. Both are of naturally occurring remindings of the sort studied by Schank in *Dynamic Memory* (Schank 1982).

Flight of the Phoenix

X was working late on a project that was due in a matter of days. As he saw the deadline approach, he considered the following two plans: either continue to work straight through the night (and the next day) or get a good night's sleep and come back to the office refreshed. The first plan allows the use of all of the time for work on the project. The second provides less time, but the time it does provide is of a better quality.

While thinking about his problem, X was reminded of a scene from the movie "Flight of the Phoenix" in which a character played by Jimmy Stewart had to start a damaged plane in the desert. The plane's ignition used explosive cartridges and Stewart had only seven left. The plane's exhaust tubes were filled with sand—which could be blasted out using the same explosive cartridges.

Jimmy Stewart's character was faced with a choice: either try to start the plane using the cartridges directly, or use some of them to clear the exhaust tubes thus enhancing the overall utility of the other cartridges.

This is a classic problem-solving situation involving resource conflicts. The reminding is of a problem analogically similar to the problem facing X, and it carries with it information about how to make the decision. In the next example, the reminding is much more concretely related to the task at hand.

Missed Exit

Y was driving along an expressway in the left lane, because traffic was moving faster there. He spotted the exit where he wanted to get off. Unfortunately, by the time he worked his way into the right-hand lane, he had missed the exit.

The next time Y was driving on the same expressway in the left lane, intending to get off at the same exit, he remembered the exit earlier than before, but still too late to make the exit.

The third time, Y made the exit.

This example, due to Phil Agre (1990, personal communication), demonstrates the role that remindings can play in repeated situated activity. By remembering earlier failures, Y is able to modify his behavior incrementally until he reaches an optimal solution for balancing the conflicting subgoals of driving fast and catching the exit.

In addition to their relevance to the abstract versus surface feature dispute, these two examples represent opposite ends of the planning spectrum in AI. Together, they challenge us to devise a coherent model of memory that supports both behaviors.

Analysis

Although we use the words "remember" and "reminding" to describe the actions of memory in each of the above examples, the dynamics which give rise to the remindings are very different, and the remindings themselves play different functional roles. These differences have important implications for the features used for memory indexing and the mechanisms by which memory functions.

In the Flight of the Phoenix example, X faces a choice between two plans to satisfy a single goal. The reminding from the movie comes about as a by-product of contemplating the two choices, and uncovering the nature of the sleep-then-work plan. The situation faced by Jimmy Stewart shares no surface features with the write-or-sleep decision. On the level of abstract problem-solving, however, the two situations match exactly. Each protagonist has a goal and a limited resource. The protagonists must decide between two plans for using the resource; one plan uses the resource directly in service of the goal, while the other splits the resource between an optimization step and a direct-use step.

The reminding plays a valuable functional role in that it gives X a different perspective on the problem at hand. In the movie, Jimmy Stewart went with the clear-and-start plan over the objections of his companions, but from X's point of view the important information comes from the debate over whether the charges that remained after clearing the exhaust would be sufficient to start the plane. This debate focuses X's attention on the question of whether sleeping would leave enough time to complete the project.

In the Missed Exit example, remindings play a different functional role, contributing to a change in behavior within a repeated activity. Structurally, the problem is to optimize over two goals: drive fast (hence drive in the left lane), and get off at the right place. However, Y can act only on the basis of immediate environmental cues and memory.

The classical planning approach to this problem, the kind of approach used successfully in the Flight of the Phoenix example, would require an uncovering of

the deep structure of the problem. However, the structure is minimal in this case; a structure-level "solution" would be look like "Move into the right-hand lane only when you need to," which is clearly no solution at all.

Failing a solution on the abstract level, the classical planning approach would use projection. For projection, a planner would require knowledge of the speed of traffic in each lane. It would require knowledge of the traffic density, since that affects how quickly one can change lanes. It would require exact knowledge of distances between landmarks.

Yet distance between landmarks is not commonly known with any precision, and traffic changes chaotically, making precise measurements and predictions impossible. A classical planner cannot hope to decide *a priori* where to begin changing lanes, because the necessary information just isn't available.

No useful solution to this problem can be defined in terms of the structure of the situation. The solution must necessarily be defined in terms of the surface features that define the activity. Those features are detectable only at the time of execution, and imperfectly even then, rendering projection impossible.

Runtime trial-and-error based on immediately perceptible environmental cues is the only practical way to solve this problem. We need to remember where we began our move last time, what happened, and how traffic today compares to traffic last time. That kind of information is exactly what the remindings in this example provide.

Implications for Memory

The Flight of the Phoenix and the Missed Exit examples tell us that memory responds to the functional needs of the situation. The functional needs of the two examples place different vocabulary and structural demands on memory. This section defines what those demands are.

As he remembers "Flight of the Phoenix," X is able to recall many surface-level details. These include the leading actor, the setting, the plane, the means for starting the plane, the two plans for starting the plane. From this we draw the uncontroversial conclusion that the vocabulary of memory includes the common semantic vocabulary of people, planes, and deserts.

Additionally, X remembered more complex aspects of the particular scene. He remembered the rationale for the clear-and-start plan and the concern that, after clearing the exhausts, the remaining cartridges would be insufficient. Thus memory includes a vocabulary associated with planning, including notions of resource limitations and of optimization.

We should point out that deep knowledge of the situation can reduce the overhead in storing some of the details. For example, the relationship between the setting (the desert) and the problem (sand in the exhaust) makes it easy to reconstruct either detail if one can remember the other.

By what features is the scene from "Flight of the Phoenix" indexed? X could probably recall that movie in the context of movies starring Jimmy Stewart, or set in the desert, or involving damaged planes. However, what makes the Flight of the Phoenix example interesting is that X's current dilemma (whether to write or sleep) contains none of these surface details, but matches the movie perfectly on the abstract planning level. We can only conclude that the movie scene is also indexed by and retrieved on the basis of those abstract planning characteristics.

Unfortunately, empirical psychologists have been unsuccessful at obtaining deep-structure remindings in a controlled setting (Gentner and Landers 1985, Ratterman and Gentner 1987, Ross 1987, Gick and Holyoak 1980, 1983, Holyoak and Thagard 1985, Holyoak and Koh 1987). Their experience is that memory retrieval occurs primarily on the basis of surface features; retrieval on the basis of abstract structure is rare and hard to reproduce.

One possible explanation for their results is that the structure and functioning of memory itself limits our ability to have such remindings. However, the Flight of the Phoenix example contradicts this explanation. We believe that memory is equally well-equipped for deep remindings and surface-feature remindings; any limitation arises from the task at hand.

We find that much of the psychologists' data comes from shallow story understanding and problem-solving tasks. In (Gentner and Landers 1985), the task was to read stories generally two paragraphs in length, and to write down any remindings the subjects had of an earlier group of stories. The stories themselves were not difficult to read, and from the subjects' point of view, *any* reminding was sufficient.

The predominance of surface-feature remindings in this context does not surprise us. Seifert et al (1986) also observed the absence of priming effects (implicit activation) during story processing on the basis of thematic structure similarity. However, when the subjects were instructed to attend to thematic structure, priming effects were discovered.

In (Gick and Holyoak 1983), subjects were given a story, referred to as "The General," in the context of a recall task. Later they were presented with a problem to solve, known as the "radiation problem."

In the story, a general wishes to overthrow a dictator. The dictator occupies a fortress in the country.

Several roads lead into the fortress; however, all of them have been mined such that while small groups can pass safely, any large force will detonate the mines. The general's solution is to partition his army among the roads, so that each group is small enough to pass safely, yet when they converge on the fortress they will be strong enough to capture it.

The "radiation problem" posits an inoperable tumor which must be killed with radiation. Any radiation strong enough to damage the tumor will also damage the surrounding tissue. The problem is to destroy the tumor with radiation while preserving the surrounding tissue; the solution is to radiated the tumor from several directions with rays weak enough that they do not damage the healthy tissue. Converging on the tumor, however, they sum together to kill the tumor.

The "radiation problem" bears little surface similarity to the story of the general, yet it was intended to be structurally analogous. Subjects found it difficult enough that few of them were able to solve it without a hint directing them to use the solution in the story.

Our problem with this study is that although the solution to the radiation problem and the solution used by the general are beautifully analogous, the *problems* are not. The goal in the story is to upset the status quo. The goal in the problem is to remove an invasive entity to *restore* the status quo. Seifert and Gray (1990) changed the General story so that the fortress is occupied by terrorists instead of by the dictator, spontaneous analogical recall was doubled.

In their conclusion, Seifert and Gray noted that the ability to be reminded based on abstract features requires that both episodes be encoded the same way. The abstract features must be available at both storage and retrieval times. What makes the Flight of the Phoenix reminding possible is that the structure of the situation is explicitly stated in the movie, allowing that structure to be used as an storage index, and the structure is uncovered by X as he ponders his dilemma, allowing it to be used as a retrieval index. Simple processing tasks will not evoke an encoding sufficiently rich for deep-structure analogical retrieval.

The remindings in the Missed Exit example key off of surface features, and we have seen that there are good functional reasons for that. What is peculiar about this example, however, is the change in behavior over time. How does the reminding occur farther back along the expressway each time?

As explained above, the circumstances are too complicated to determine an exact point at which to begin changing lanes. This is true *a posteriori* as well as *a priori*, and for the same reasons. We hypothesize, therefore, that when the failure occurs it gets indexed by associated landmarks. "Associated landmarks" would include landmarks near the exit, and probably landmarks near where Y began to change lanes. A conscious version of this indexing would be something like "The bridge wasn't soon enough. Next time I need to start moving earlier."

How does this help next time? We note that a good driver does not focus his attention directly in front of her car, but scans the road far ahead. In situations where his view of a familiar road is blocked, as around a corner or in fog, the good driver does a small amount of projection, anticipating what will come next. Thus, before he gets there, Y will either see the landmark with which the failure is associated, or will imagine it. Seeing it or imagining it is enough to trigger the reminding. If the reminding occurs far enough back along the expressway, Y will catch the exit. Otherwise, he has another episode to remember next time, indexed by a landmark farther back.

More Examples, More Functionalities

The people involved in the Flight of the Phoenix and the Missed Exit examples had different functional requirements: X had to decide between two alternative courses of action; Y had to remember to begin moving toward the right-hand line earlier. Those different functional requirements caused caused memory to behave in different ways. Next we consider two examples which seem to place the same functional requirements on memory, but yield different behavior. These behaviors reveal additional subtleties in the organization of memory.

Milk

J ran out of milk during breakfast. He couldn't buy any then because he had to go to work.

While running an errand later that day, J passed a grocery store. He remembered that he needed milk. Deciding that the milk would keep until he went home, he stepped into the store.

J has a goal to get milk. He remembers the goal at the best possible moment for satisfying it, when the conditions in the world are ripe for buying milk.

The next example shows that we can't always recognize even the most obvious of opportunities.

Steve's Loan

Steve loaned his good friend Tom $20 one Friday night. The next morning, Tom found the $20 bill, unspent, and told himself to give it to Steve when he saw him next. Later, Tom went to Steve's apartment to watch a ball game, but didn't

remember the money until Steve happened to take some dollar bills out of his pocket.

Tom knew in advance the conditions of a good opportunity to give Steve the money: seeing Steve. For some reason the effort to remember at the right time failed, however, and he remembered thanks only to an unexpected event.

Analysis

The remindings in these two stories share the common theme of remembering to do something. The interesting difference between them is that, in the first case, the reminding came at a "good" time; in the second, the reminding was triggered by a causally unrelated and unpredictable event.

Running out of milk gave J the goal to get milk, but he couldn't couldn't pursue the goal immediately. What happened to that goal? We could imagine him forgetting about it, not remembering that he was out of milk until the next time he needed some. Or, we could imagine him planning to buy milk at a specific time of day, such as during lunch, or perhaps anticipating the errand. The former is an unfortunate lapse which happens to all of us sometimes; the latter requires projection, which has been proved intractable in the general case (Chapman 1985).

In this case, however, J made no effort to plan how a trip to the grocery store might fit into the tumultuous events of the day, nor was the goal to get milk forgotten. Instead, it remained in memory indexed by features which would indicate a good opportunity to buy milk, i.e. the image of a grocery store. This is an example of *opportunistic memory*, an idea examined in (Converse, Hammond, and Martin, 1990).

According to the opportunistic memory theory, goals and plans are governed by *policies* (McDermott 1978). A policy describes a state of the world which, while desirable to an agent, does not require constant attention. Having milk is a typical policy. We don't spend our days asking if we have milk, but if we run out, we notice it and make some effort to get more.

Well-worn goals like obtaining milk have well-worn plans indexed under them. In opportunistic memory, when the agent decides that he cannot immediately pursue the goal, the goal is *suspended*, or saved in memory. The preconditions of the plans for satisfying it are used as indices, so that if the world presents a situation in which many of the preconditions are satisfied, the agent can take advantage of the opportunity.

Getting milk is a common task which J has accomplished many times, and almost always at a grocery store. The precondition of being at a store are well-known and easily recognized. Thus when the goal of getting milk is suspended, J knows under which the features to index it. When the opportunity arises to buy milk, J can recognize it and take advantage of it.

The example of Steve's Loan provides an interesting counterpoint to the first because it seems to betray the opportunistic memory theory. Tom had a goal to return the money. He made an explicit, conscious attempt to index that goal under the conditions that would make it easy to satisfy: seeing Steve. Yet, the goal was not triggered when those conditions were met, but only on the basis of surface features which could not have been anticipated.

The trouble arises out of a need to avoid spurious remindings. Some features are too common to be used as reliable indices, unless there is a close two-way correlation between the feature and what is to be remembered. For Tom, the sight of Steve is associated with too many things to be a good feature on which to index anything. They commute to work, cook dinner, play pool, and watch TV together. Grocery stores, on the other hand, have a unique association with purchasing food. Although the sight of a grocery store is common, the role grocery stores fill in our lives is narrowly enough defined that they make a good feature for remindings about purchasing food.

Simple Remindings

In contrast to the complex, structure-rich reminding in the Flight of the Phoenix example, many, probably most of our remindings are extremely simple.

Hawks

X says to Y, "I saw a hawk yesterday." Y says, "Oh, I was just reading about hawks. Do you know what kind it was?"

Here is a surface-level reminding with no apparent functionality. Consider the task, however: Y has only to understand what X has said to him and react to it in an appropriate way. The subject of the utterance is a hawk-sighting; the context of the utterance is informal conversation about recent personal events in one's life. Hence the reminding is doubly appropriate. In the process of relating the utterance to his relevant knowledge of hawks, Y is reminded of a recent activity suitable for conversation. Simple remindings of this kind are a part of basic human understanding.

Summary

This paper presented five reminding examples. First, a problem-solving situation generated an reminding which, while sharing no important surface features, provided insight into the nature of the problem. In the second example, remindings occurred in the midst of repeated, situated activity on the basis of perceptually evident features. The remindings helped the agent avoid failure and optimize over conflicting subgoals. Third, a goal that could not be pursued immediately was suspended and indexed in memory by the features that would indicate a good opportunity for its satisfaction. The fourth example revealed a limitation of this functionality: goals cannot be indexed by features which are extremely common and associated with many other goals. Finally, informal conversation provoked a simple reminding nonetheless well-suited to the task of understanding an utterance and generating an appropriate reply.

These five examples reveal the diversity of situations in which remindings occur, and the diverse roles they play. Remindings help us solve problems, they help us optimize between conflicting goals, they help us notice opportunities, they help us do what we intend to do, and they make it possible for us to relate new information to what we already know.

In response to the debate over whether remindings occur on the basis of surface features or deep structure, we have argued that the context and type of activity is the important factor, rather than any inherent limitation of memory. Tasks requiring abstract analysis yield abstract remindings; tasks that depend on features in the world index on the basis of readily-identifiable surface features. Finally, we have argued that this influence of activity on memory is functionally significant: abstract remindings are more likely to carry the kind of information we need to solve complex problems; surface-level remindings are more likely to help us react quickly to our environment.

References

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

Birnbaum, L, & Collins, G. (1989). Remindings and Engineering Design Themes: A Case Study in Indexing Vocabulary. In Proceedings of the Second DARPA Workshop on Case-based Reasoning in AI. Pensacola Beach, FL. DARPA.

Chapman, D. (1985) Nonlinear Planning: A Rigorous Reconstruction. In Proceedings of the Ninth International Joint Conference on Artificial Intelligence, pp.1022-24. IJCAI, 1985.

Converse, T., Hammond, K., & Martin, C. (1990). Integrating Planning and Acting in a Case-based Framework. In Proceedings of the Eighth Nation Conference on Artificial Intelligence, pp.292-297. AAAI, 1990.

Gentner, D. & Landers, R. (1985). Analogical Reminding: A Good Match is Hard to Find. In Proceedings of the International Conference on Systems, Man and Cybernetics, Tucson, AZ.

Gick, M. & Holyoak, K. (1983). Schema Induction and Analogical Transfer. *Cognitive Psychology,* 15:1-38.

Hammond, K. (1989). *Case-based Planning: Viewing Planning as a Memory Task.* Academic Press.

Holyoak, K. & Thagaard, P. (1985). Discovering the Wave Theory of Sound: Inductive Inference in the Context of Problem Solving. In Proceedings of the Ninth International Joint Conference on Artificial Intelligence, pp 610-612, Los Angeles, Ca. AAAI.

Kolodner,J. & Simpson, R. (1989). The Mediator: Analysis of an Early Case-based Problem Solver. *Cognitive Science Journal.*

McDermott, D. (1978). Planning and Acting. *Cognitive Science,* 2, 1978.

Rattermann, M. & Gentner, D. (1987). Analogy and similarity: Determinants of accessibility and inferential soundness. In Proceedings of the 9th Annual Conference of the Cognitive Science Society. Hillsdale, NJ: Lawrence Erlbaum Associates.

Ross, B. H. (1987). This is Like That: The Use of Earlier Problems and the Separation of Similarity Effects. *Journal of Experimental Psychology: Learning, Memory and Cognition,* 13(4):629-639.

Schank, R. (1982). *Dynamic Memory: A theory of reminding and learning in computers and people.* Cambridge University Press.

Seifert, C. & Gray, K. (1990). Representational Issues in Analogical Transfer. In Proceedings of the Twelfth Annual Conference of the Cognitive Science Society.

Seifert, C., McKoon, G., Abelson, R.P., & Ratcliff, R. (1986). Memory connections between thematically similar episodes. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* 12, 220-231.