

Reconstructive Memory: A Computer Model*

JANET L. KOLODNER

Georgia Institute of Technology

This study presents a process model of very long-term episodic memory. The process presented is a reconstructive process. The process involves application of three kinds of reconstructive strategies—component-to-context instantiation strategies, component-instantiation strategies, and context-to-context instantiation strategies. The first is used to direct search to appropriate conceptual categories in memory. The other two are used to direct search within the chosen conceptual category. A fourth type of strategy, called executive search strategies, guide search for concepts related to the one targeted for retrieval.

A conceptual memory organization implied by human reconstructive memory is presented along with examples which motivate it. A basic retrieval algorithm is presented for traversing that structure. Retrieval strategies arise from failures in that algorithm. The memory organization and retrieval processes are implemented in a computer program called CYRUS which stores events in the lives of former Secretaries of State Cyrus Vance and Edmund Muskie and answers questions posed in English concerning that information. Examples which motivate the process model are drawn from protocols of human memory search. Examples of CYRUS's behavior demonstrate the implemented process model. Conclusions are drawn concerning retrieval failures and the relationship of episodic and semantic memory.

1. INTRODUCTION

Q: Have you been to Saudi Arabia recently?

A: Yes, most recently last month, to discuss the Camp David Accords with King Khalad and Prince Fahd.

*This work was supported in part by the Advanced Research Projects Agency of the Department of Defense and monitored under the Office of Naval Research under contract N00014-75-C-1111 while the author was in residence at Yale University. It is currently supported in part by the National Science Foundation under Grant No. IST-8116892.

I am indebted to Roger Schank, who helped guide this research when I was a graduate student, to all the members of the Yale AI project, who provided an environment conducive to good research, and to Gordon Bower, Mike Williams, and Larry Barsalou, who helped me to understand the psychological implications of this work and to make this paper more intelligible.

Q: Where did you go afterward?

A: To Syria. I was touring the Middle East talking to each of the Arab leaders about the Accords.

Q: How many times have you been to the Middle East in the past six months?

A: I was in Israel and Egypt this past summer on two separate trips, and after the Camp David Summit, I was in the Middle East to talk to Arab leaders about the Camp David Accords.

Taking part in a dialog or discussion often requires retrieval of past events from memory. The answers given in this hypothetical dialog would have been reasonable ones for Cyrus Vance to give while he was secretary of state. Yet, he went on many trips as secretary of state and many more during the rest of his life. How can the appropriate events for answering questions be extracted from memory? What are the processes that allow this retrieval to happen?

The problem of searching a large memory efficiently is one of the biggest problems facing designers of large systems, yet one which has been sorely neglected. While there has been a lot of work done on problems of knowledge representation and inference processes, the assumption or simplifying condition of most investigations has been that necessary information (e.g., frames, scripts) would be available when needed.

One way to approach the problem of searching a large memory is to look to people. In a normal day, people effortlessly recall past events and episodes from their lives many times. If we knew how people did it, then we could model our large systems after people. Human memory is often described as being reconstructive (Norman & Bobrow, 1977; Schank, 1980; Spiro, 1979; Williams & Hollan, 1981). Psychologists have been describing reconstructive memory as a problem-solving method—a way of putting together little bits and pieces to come out with a whole (Williams & Hollan, 1981). Within artificial intelligence, we have also been looking at reconstructive memory, though at different problems and at a different level of detail. Because we want to implement our theories on the computer, we must be able to explicitly state particular processes involved in retrieval and the memory organization which supports those processes. This study presents memory organization and retrieval strategies for reconstructive memory, along with a computer program which implements the theory.

CYRUS (Computerized Yale Retrieval and Updating System; Kolodner, 1978, 1981a, 1984; Schank & Kolodner, 1979) is a computer program which implements the theory of long-term memory presented here. CYRUS is meant to be both an intelligent fact retrieval system and a model of long-term memory for events. CYRUS has been designed to keep track of events in the lives of important people. It answers questions posed to it in English

pertaining to that information. Thus, it both organizes events in its memory and searches memory reconstructively to retrieve them.

The following protocol is of CYRUS answering a number of questions. Notice in the intermediate output (between the question and the answer) that the reasoning that CYRUS goes through in getting its answers is very similar to, or at least would be reasonable for, what a person would do in answering the same questions.

@CYRUS

Are you interested in Muskie or Vance? (M or V): *VANCE

*(DIALOG2)

>When was the last time you were in Egypt?

inferring a diplomatic trip

answering question using time context

ON DEC 10 1978.

>Why did you go there?

answering question using previous context

TO NEGOTIATE THE CAMP DAVID ACCORDS.

>Who did you talk to there?

answering question using previous context

inferring undifferentiated political meeting

WITH ANWAR SADAT.

>Has your wife ever met Mrs. Begin?

inferring a social political occasion

YES, MOST RECENTLY AT A STATE DINNER IN ISRAEL IN

JAN 1980.

>What heads of state have you met?

inferring a diplomatic meeting

searching directly for \$MEET

...

applying strategies to search memory

...

searching for sM-CONFERENCE

...

searching for sM-VIPVISIT

searching sM-VIPVISIT instance for \$MEET

...

searching for I-NEGOTIATE

...

KING HASSAN, MENACHIM BEGIN, MUSTAFA KHALIL, IAN SMITH, PRESIDENT ASSAD, KING HUSSEIN, AND ANWAR SADAT.

>Last time you went to Saudi Arabia, where did you stay?

inferring a diplomatic trip

answering question using time context

IN A PALACE IN SAUDI ARABIA ON SEP 23 1978.

>Did you go sightseeing there?

answering question using previous context
YES, AT AN OILFIELD IN DHARAN ON SEP 23 1978.

In working on CYRUS, we had two objectives. One was to enable the computer to do the kinds of intelligent memory tasks that people can do—remembering. The other objective was to use the computer as a tool to investigate human reconstructive memory. In exploring reconstructive memory, there are three things to look at:

1. the memory organization
2. the processes for retrieving information and facts from that organization
3. the maintenance of that organization over time as additional items or experiences are added to memory.

Here we will look at the memory organization and the retrieval strategies. To find out more about memory maintenance over time, see Kolodner (1983). For more detail about all three of these topics, see Kolodner (1984).

1.1 What Is Reconstruction?

For the most part, people are expert at recalling information about their pasts. When people try to retrieve information about a particular episode in their lives, however, they often find they do not have the entire episode immediately available. In that case, they attempt to *reconstruct* the rest of the episode. When asked to recall a recent trip one person started his reconstruction as follows:

Let me see. When we got there, we must have had to find a hotel. Yes, I remember, we had a guide book with a lot of hotels in it, and we called a few of them until we found one with a vacancy. It was late afternoon, so we must have gone out to eat soon after that.

Clearly, the trip was not stored in memory in one large chunk. Instead, it must have been stored in bits and pieces requiring reconstruction to put them back together. Such reconstruction can happen through application of generalized knowledge. In the protocol quoted, the knowledge that “one must have a hotel to stay in during a trip” and “usually after arrival, the next step of a trip is to check in at the hotel” allowed “finding a hotel” to be recalled.

The process seems to involve building a description of “what must have happened” and then filling it out with “actual” details. We can thus conclude the following about human memory for events:

Memory Principle #1

Human remembering is often a process of reconstructing what must have happened rather than directly retrieving what did happen.

This process of *reconstruction* is almost always used by people trying to remember episodes, as evidenced by the fact that people tend to “remember” details incorrectly so often. In fact, retrieval confusions and false recall must be inherent to a process which produces probable rather than actual explanations. We will see later why we might want to give this capability to the computer, despite the fact that it often produces errors in details.

Within psychology, approaches to reconstructive memory have stressed a theory of why it is done (Williams & Hollan, 1981), experiments that show it is done (Spiro, 1979; Williams & Hollan, 1981) and some of the strategies employed (Williams, 1978). They have purposely avoided describing a memory organization or basic retrieval process (Williams & Hollan, 1981). An AI approach to the problem must stress these two problems in addition to describing the reconstructive retrieval strategies. Without these two components, we cannot implement and test our theory. Thus, as a prerequisite to explaining reconstructive strategies themselves, we must explain memory organization and access. Retrieval strategies stem from failures in the basic retrieval process.

Before continuing, we must note an important difference between our use of the word “reconstruction” and the usual use of the word in psychology. Psychologists use reconstructive memory to refer to the producing of plausible scenes or making plausible guesses at what a targeted event might have been. This model, on the other hand, uses reconstruction to refer to constructing a description of a target event, adding features to progressively differentiate it from its nearest neighbors, and finally finding its hiding place in a well-organized memory. There are two primary differences in emphasis. First are assumptions about the organization of memory. To psychologists, organization of items in memory does not come into play in reconstruction. In this model, memory’s organization drives the process. The second difference in emphasis concerns the process itself. We have, in effect, included a verification process in the reconstructive process. Thus, reconstruction is a two-step loop: (1) coming up with a description, and (2) discovering if it is in memory, or searching for it, and going back to (1). In this way, the contents of memory drive the reconstructive process. Because verification is part of the procedure, however, mistakes are not always obvious. While errorful descriptions of a targeted event may be produced, the verification process or search through memory suppresses their output. The primary problem with this is that it does not directly explain how people can generate gross errors and yet have complete confidence that they are correct. On the other hand, the search (or verification) part of the process allows memory to guide the generation of hypothetical details of a target event.

We must also state the goals of our study at this point. Our major goal is to look at reconstruction of day-to-day events. Our bias in this direction stems from observations that it is day-to-day events, and the generalizations

that can be made from them, that form the core of our ability to understand language. Episodic structures such as scripts, goals, and themes (Schank & Abelson, 1977) have already been shown to play a major role both in understanding and in getting around in the world. Such episodic structures can only arise by keeping track of the events they are derived from in an episodic memory. In keeping track of events, we must figure out how to organize them. The reconstructive abilities of people give us pointers into such an organization. We thus study reconstructive memory rather than rote memory, and episodic memory rather than semantic or factual memory.

In implementing our model, we are most interested in those parts of memory which directly contribute to reconstruction. Our model of Cyrus Vance thus does not model all of Vance's memory, but only those things necessary for reconstruction of day-to-day events during his tenure as secretary of state. Thus, we do not store what he knows about Egyptian history or about international relations before he became secretary of state, though in a more complete memory, those things might be relevant and useful to have. Nor do we store childhood memories. The Vance memories we store are strictly those parts relevant to a restricted interpretation of his secretary of state role.

Another restriction on our model is that we are not dealing with recency or frequency data. While some questions we use as examples seem to depend on such knowledge, it is only in CYRUS in the most rudimentary way. What we are really interested in in those questions is the capability of finding anything at all in a large store of events. Thus, in answering "how many," our intention is to show how a large number of events can be found without enumeration, rather than showing that CYRUS makes frequency judgements.

The remainder of this study is divided into four parts. First a memory organization to support reconstructive retrieval is presented. Second, a basic retrieval process for traversing that organization is discussed. Third, problems with that process are enumerated, along with reconstructive retrieval strategies which derive from those problems. In the last section the implications of the processes proposed are discussed. In motivating the problem of memory organization, we present some of the capabilities and failures of human reconstructive memory.

2. MEMORY ORGANIZATION FOR RECONSTRUCTIVE RETRIEVAL

In discovering a memory organization for reconstructive memory, we must find one that not only supports such retrieval, but also *requires* it. To begin addressing that issue, consider again the following hypothetical dialog given in the introduction:

Q: Have you been to Saudi Arabia recently?

A: Yes, most recently last month, to discuss the Camp David Accords with King Khaled and Prince Fahd.

Q: Where did you go afterward?

A: To Syria. I was touring the Middle East talking to each of the Arab leaders about the Accords.

The answers given in this dialog would have been reasonable ones for Cyrus Vance to give while he was secretary of state. During the four years he was in office, however, he went on hundreds of trips, among many other memorable activities. How could he choose the correct ones from memory to answer specific questions? If memory were arranged in lists, then we could imagine a process which searched down an appropriate list until the correct episode was found. This seems unreasonable, however, for a memory with hundreds or thousands of events of the same type: Searching down a long list is a slow process which becomes slower as additional items are added to the list. People's memories do not slow down as they learn more (Smith, Adams, & Schorr, 1978).¹

Furthermore, if people searched down long lists in their memories to retrieve events, then it would be easy for them to enumerate experiences of particular types. For instance, a question such as the following would be easy to answer:

(Q1): Recall all the times you have been to museums.

This question, however, is not easy for people to answer, as suggested by the following protocol:

I know I've been to a lot of museums in Europe. I've been to England, and I went to a number of museums there—some in London—the British Museum, the National Gallery, and a few smaller galleries. . . . I was at a museum in Brighton—the Royal Pavilion. I've been to museums in Paris—the Louvre and some smaller ones. In Rome, I've been to. . . . In Naples, to. . . . In Florence, to. . . .

The lists of experiences people enumerate are constructed on the fly as they are answering a question. In the protocol cited, the person tried to recall "experiences in museums," "experiences in museums in Europe," "experiences in museums in England," "experiences in museums in Lon-

¹In fact, there is some debate about this among psychologists. Anderson (1974, 1976) cites the "fan effect" as evidence that retrieval slows down with the addition of new items about a particular concept. Smith et al. (1978), on the other hand, has shown that retrieval does not slow down with the addition of new items when context is guiding the retrieval. Reder and Anderson (1980) conclude that when people make consistency judgments, rather than retrieving actual facts from memory, there is no interference to slow down the retrieval process.

don," etc., filling in (or reconstructing) additional details with each iteration. Generalizing from that, we can make the following hypothesis:

Memory Principle #2

Remembering requires progressive narrowing in on a description of the event to be remembered.

We can describe that process as a process of specifying details which can differentiate the targeted event from other events in memory.

The protocol above and others like it suggest that narrowing the description of a sought event, or focusing in on an event by providing more detail is a *necessary* part of remembering (Williams, 1978). Thus, we can state the following principle of memory retrieval:

Memory Principle #3

In a reconstructive memory, memories are not directly enumerable. Instead, the features that describe individual memories must be reconstructed.²

What kind of memory organization would make such a retrieval process necessary? One thing we know about human memory is that it is organized around concepts. Since a "visit to the National Gallery" and a "visit to the Louvre" are closely related concepts, we can imagine that they would be reasonably organized in memory in close proximity to each other. Since both are "visits to museums," both must be related to that concept. Memory organization can reflect this proximity in a way similar to that illustrated in Fig. 1.

In this organizational structure, similar concepts are placed in close proximity to each other in a network. The links are bidirectional and form an inheritance or ISA hierarchy. Each concept points back (up) to the concepts it is an instance of (through ISA links), and forward (down) to all of its example concepts (through INSTANCE-OF links). Nodes in the network can be described in terms of the more general concepts they are instances of. Using this network, we can describe the "visit to the Louvre" as a "visit to a museum in Paris" and a "visit to an art museum." The graph can also be traversed downward. Starting from "visits to museums in Europe," "visits to museums in London" and "visits to museums in Paris" can be retrieved.

²Of course, there must be a capability for doing some types of enumeration. People remember short lists, for example, the twelve months of the year. Our thesis is that these lists are either committed to memory or are derived from a well-structured concept which serves as a map or chart. People report generating the list of 50 states, for example, by mentally walking around a map of the United States.

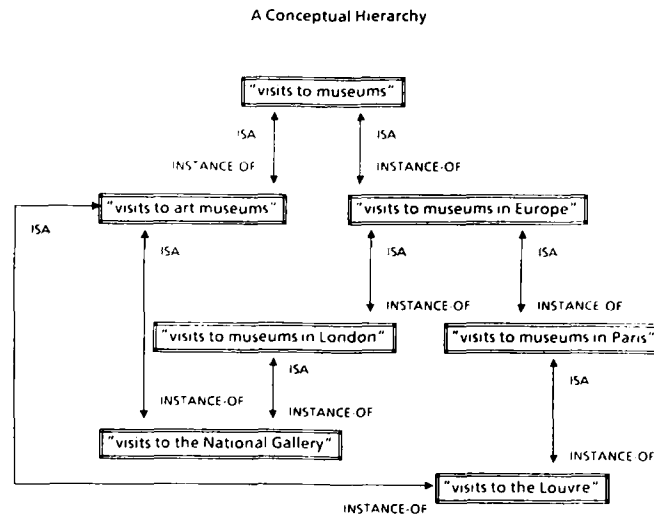


Figure 1

Although this network both organizes similar concepts in proximity to each other and allows traversal from more general to more specific concepts, it has a major flaw: it allows too much. In an organization such as this, there would never be any possibility of forgetting, a phenomenon frequently observed in humans. Nor does it explain false starts such as the following:

Was I in a museum in Oxford? No, I was only there for a couple hours, and I only had time to see the outsides of the buildings.

In this fragment, the respondent has attempted a reconstruction of a possible museum visit by considering a "museum visit in Oxford." However, after further reasoning, he decides that this event never happened. In a fully enumerable memory, this phenomenon would not occur. Reconstructive specification of details would not be necessary, if arcs were all labeled by ISA and INSTANCE-OF. All examples could be retrieved and checked without guessing at details. No contrafactual episodes would be considered and no false starts would be observed.³

While we do see false starts and forgetting in people, it may seem strange to require such behavior from a computer memory. Computers, after all, have large memories and never need to forget. Or do they? The

³Note that if we had a strength associated with each arc in the network above and thresholds for cut-off or output interference, then we would get a memory that forgets and that cannot enumerate. While reconstructive retrieval strategies could be laid on top of such a structure, the structure itself would not provide enough information to guide and control reconstruction. We shall see that the memory structure to be proposed can actually aid the reconstructive process.

major problem in a computer program with a memory organization like the one in Fig. 1 would be control of memory search. If all instances of a concept were related through INSTANCE-OF links, then all examples of any concept could always be enumerated. While this would not be a problem if a concept had only a few instances, consider the implications of this organization in a memory where each concept had hundreds or thousands of instances. To find any particular instance, all instances would have to be enumerated, and the correct one chosen from those. This, in fact, is equivalent to sequential search of a list which we have already said is undesirable: Memory slows down with each new addition.

The major reason for problems in the structure illustrated above is that the relations represented by its arcs are inadequate. In order to control search, the relationships in a conceptual network must be more sophisticated than INSTANCE-OF and ISA links. We can figure out what the relations represented by the arcs should be by looking again at the observed retrieval process. Reconstructive retrieval is a process of building up a plausible description of the target item. One step of the process involves narrowing a broad concept that describes too many similar items by specifying additional features. After the concept is narrowed to describe only one item, that item can be found in memory. The arcs, then, should represent features of items organized by a particular concept. Specification of one of these features will allow traversal of the corresponding arc, arriving in memory at a more specific concept. Part of the memory structure implied by the protocols above is exhibited in Fig. 2.

The arcs are labeled by features of the concepts they point to, and thereby serve to *differentiate* those items from other similar items. Thus the observed retrieval process implies a memory organization with the following feature:

Memory Principle #4

Similar items in memory are organized around the same concept according to their differentiating features.

We must put two additional restrictions on this organization. First, its arcs must not be enumerable. If they were, then so would all of the associated nodes. This restriction also explains the false-start fragment quoted previously where the person was wondering if he had been to a museum in Oxford. If arcs are not enumerable, then their labels must be reconstructed during retrieval. Reconstruction can produce possibilities which are not in memory.

The second restriction follows from the first. Arcs may be traversed only if their labels have been specified. Such an organization will prohibit enumeration, but will allow directed traversal through memory. Intuitively,

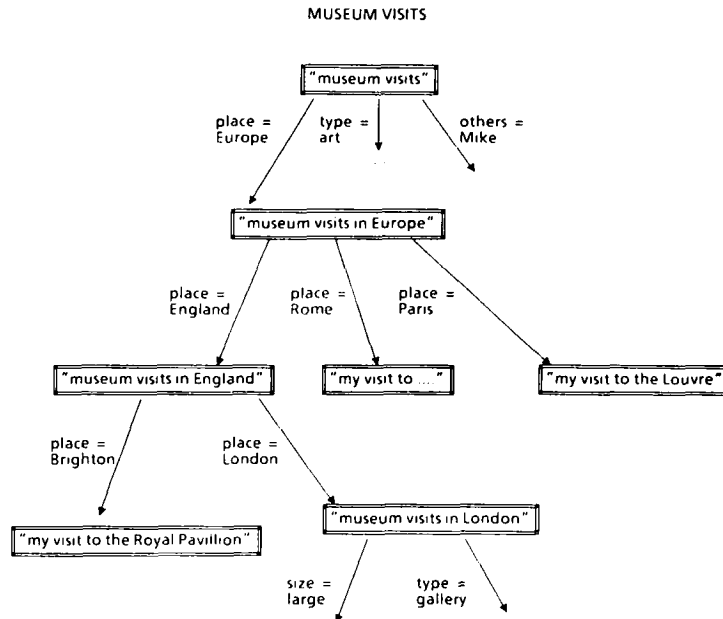


Figure 2

we can think of these structures as “locked conceptual networks.” The arcs can be thought of as locks which can only be opened (and thus traversed), if the correct key is available to open them (i.e., if the label on the arc can be specified). In this way each concept organizes all of the concepts related to it, but obtaining any of these related concepts requires specifying the feature designated as its path “lock” or key.

Such an organization accounts for some of the reconstructive memory phenomena we observe in people. But why use such a restrictive memory organization in a computer system? While it may seem that a locked network requires more work for retrieval, the alternatives require enumerating, and therefore examining, many more items. In a large memory, directed retrieval is preferable to enumeration. A locked memory structure is oriented toward a directed retrieval process involving incremental specification of possible features of items to be retrieved. Instead of examining a large number of items, extra effort is made to direct search to a small number of the most relevant items. Reconstruction is a method of directing search to only relevant items in memory.

2.1 E-MOPs

The concepts which organize *events* in memory will be referred to as Episodic Memory Organization Packets (E-MOPs). Labels of arcs in an

E-MOP are called its *indices*, and are said to *index* the items organized by an E-MOP (i.e., events). When an event is entered into one of these structures, the features which differentiate it from other items in the E-MOP are extracted and used to index it. To retrieve events from an E-MOP, features of a targeted event are specified, the corresponding arcs are traversed, and the concepts indexed by the specified features are retrieved.

The E-MOPs implicitly and explicitly referred to in the examples above were "museum visits," "trips," "finding a hotel," "diplomatic trips," and "diplomatic meetings." In general, we can expect an event memory to have E-MOPs for each major type of event it knows about. Thus, E-MOPs for day-to-day events might include "getting up in the morning," "eating in a restaurant," "going to the movies," "driving to the office," etc. Memory for a diplomat will also include E-MOPs for "diplomatic trips," "diplomatic meetings," "negotiations," "speeches," and other diplomatic activities. The E-MOPs necessary to organize events in the life of a particular person will depend on that person's activities and experiences. No matter what type of experiences it is organizing, however, each E-MOP will organize its events according to their differentiating features.

E-MOPs can be thought of as conceptual categories, and will often be referred to in that way. It must be pointed out however, that they are not categories in the conventional sense of the word. Although they organize similar concepts, their members cannot be enumerated.

A script (Cullingford, 1978; Schank & Abelson, 1977) is the simplest form of an E-MOP. Like scripts, E-MOPs provide conceptual categorization for events. The emphasis of E-MOPs is different, however. While script research emphasized the structure of generalized episodes and its use in understanding, the emphasis in studying E-MOPs is the organization of individual episodes with respect to each other and to appropriate generalized episodes. Scripts were thought of primarily as processing structures, while E-MOPs are looked at as both storage and processing structures.

E-MOPs organize similar events with respect to each other by indexing them according to their differences. The similarities between the events an E-MOP organizes make up the descriptions of the generalized episodes and provide the same processing information scripts provided. In fact, this organization gives memory the ability to learn generalized episodes from particular experiences. E-MOPs will be referred to generically as "MOPs."⁴

⁴These structures are related to Schank's (1980) MOPs and to Lebowitz's (1980) S-MOPs, but the concerns in defining MOPs and S-MOPs were different than those in defining E-MOPs. In particular, Schank's (1980) concern was with showing the interrelatedness of structures in memory. Thus, in his example domain of professional office visits, he described how visits to doctors, dentists, lawyers, and other professionals are similar, and how the structures they are stored in are related. My concern, on the other hand, is with the processes for

2.2 The Internal Organization of an E-MOP

Internally, an E-MOP is a network in which each node is either an E-MOP or an event. Each E-MOP has two important aspects: (a) generalized information characterizing its episodes, and (b) treelike structures that index those episodes by their differences. As for Schank's (1980) MOPs, the generalized information associated with an E-MOP is called its *content frame*. An E-MOP's content frame holds information describing its events, including their usual cast of characters, settings props, and topics, and their usual relationships to other events.

One of the E-MOPs CYRUS uses is "diplomatic meetings." Each of Vance's or Muskie's diplomatic meetings that are entered into CYRUS's memory are indexed in that E-MOP. CYRUS knows that the participants of diplomatic meetings are foreign diplomats, that their topics are international contracts, and they include discussion between the participants about the topic, that their goal is usually to resolve a disputed contract, and that they are normally instrumental to a larger negotiating episode. That is some of the information that makes up the content frame of "diplomatic meetings." Figure 3 shows some of the normative information CYRUS knows about "diplomatic meetings."

Notice that the relationships between E-MOPs cannot be stored as simple links. Rather, they must also specify correspondences between their components. Thus, CYRUS's "diplomatic meetings" MOP specifies not only that it is related to "negotiations," but that the participants in the "negotiations" it is related to are the countries the participants of the meetings represent, and that the topic of the negotiations includes the topic of the meeting.

The second important feature of an E-MOP is its indices. Events are indexed in E-MOPs by those features which differentiate them from other events in the E-MOP. An E-MOP's indices can index either individual episodes or specialized E-MOPs. When an E-MOP holds only one episode with a particular feature, the index corresponding to that feature will point to the individual episode. When two or more episodes in an E-MOP share the same feature, its corresponding index will point to a specialized sub-MOP (with the structure just described) which organizes the subset of events with that feature. In this way, MOP/sub-MOP hierarchies are formed.

retrieval of individual episodes, and the organizational requirements those processes place on memory. If my domain were professional office visits, I would be describing how particular visits are stored in memory in relation to each other and the generalized information that would allow them to be retrieved.

Lebowitz (1980) used S-MOPs to store terrorism events derived from newspaper stories. E-MOPs are more experientially oriented and hold more detailed information than his S-MOPs.

'diplomatic meetings' (\$MEET) -- generalized information

larger episodes:

- "negotiations" (I-NEGOTIATE)
 - participants = diplomats of same nationality as meeting participants, and of other countries involved in meeting topic
 - topic = generalization of meeting topic
- "diplomatic trips" (sM-VIPVISIT)
 - destination = country or area involved in meeting topic, or country of residence of meeting participant
- "conferences" (sM-CONFERENCE)
 - participants = diplomats of same nationality as meeting participants, and of other countries involved in meeting topic
 - location = location of meeting
 - topic = generalization of meeting topic

enables:

- "treaty signing" (\$TREATY)
 - sides = sides of meeting topic

sequence of events:

- participants MTRANSing to each other about meeting topic

preceding and following events:

- "diplomatic meetings" (\$MEET)
 - participants = subset of meeting participants
 - topic = aspect of meeting topic

more general E-MOPs and classifications:

- "political meetings (sM-MEETING)
 - all components correspond
- "diplomatic activities"
 - all components correspond

role fillers:

- participants: foreign dignitaries of countries involved in contract being discussed
- location: conference room in capital city of country of residence of some important participant
- topic: international contract
- duration: one to two hours

Figure 3

The process of choosing differences for indexing is described in detail in Kolodner (1983), and thus will not be detailed here. In general, there can be hundreds of differences between two events, ranging from microscopic to abstract in level of detail. Differences indexed in E-MOPs are domain-dependent and derived from specified content frame features. Thus the features CYRUS uses are related to Vance's professional goals and include occupational, diplomatic, and social features. In addition, good indexing features make predictions about other features of the events they organize.

Consider, for example, how the following two events are indexed in CYRUS' "diplomatic meetings" MOP (see Fig. 4). Both of these meetings are diplomatic meetings with foreign diplomats about international contracts. One is with Begin about the Camp David Accords (EV1), and one is

with Gromyko about SALT (EV2). These two meetings are discriminated in CYRUS's "diplomatic meetings" MOP (\$MEET) as shown in Figure 5.³ Notice that indexing in an E-MOP is two-tiered. The first tier points to the type of index (e.g., participants, topic), the second its value for a particular event or group of events (e.g., Begin, SALT). In the actual CYRUS implementation, EV1 and EV2 point to recordings of the particular events. In a more sophisticated implementation, they need only include some subset of those features of the event left over after indexing.

- EV1 "diplomatic meeting" with
actor (Vance)
others (Begin)
topic (the Camp David Accords)
- EV2 "diplomatic meeting" with
actor (Vance)
others (Gromyko)
topic (SALT)

Figure 4

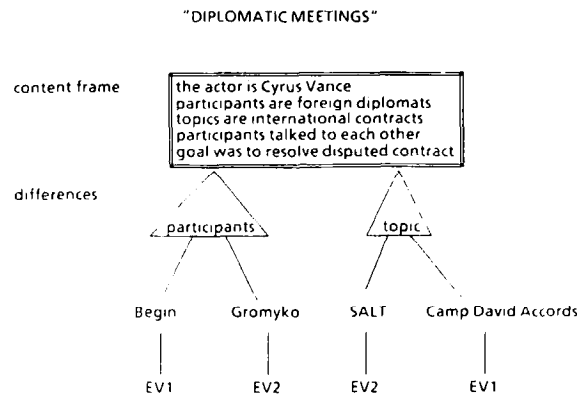


Figure 5

As additional meetings are added to memory, generalized content frame information is refined, and additional indices for events are created. As that is happening, new E-MOPs are created were multiple episodes are indexed. Each of those new E-MOPs has a content frame based on the similarities between the episodes it indexes. Episodes are indexed in each new E-MOP according to their differentiating features. These newly created specialized E-MOPs inherit content frame properties from the more general E-MOPs they are specialized from, and in addition have their own more specialized content frame information. Thus, E-MOPs and their specializations form a hierarchy discriminated by differences from content frame features.

³The content frame has been simplified in this and subsequent drawings.

After many meetings with Begin are added to the memory structure in Fig. 5, its organization would include the arcs and nodes in Fig. 6. MOP1, "diplomatic meetings," is a refined version of the diplomatic meetings MOP in Fig. 5, while MOP2 and MOP3—"diplomatic meetings with Begin" and "diplomatic meetings about the Camp David Accords," respectively—are at the points in MOP1 where the meeting with Begin about the Camp David Accords is indexed. SALT and Gromyko remain indices to EV2, an individual event, since no additional similar meetings were added to the MOP. Thus there would be no MOPs created at those index points until additional meetings about SALT or with Gromyko were added. Index points (2) and (3) index meetings with Dayan and about Jerusalem, respectively, in "diplomatic meetings." Index points (6), (7), (8), and (9) are new indices in MOP2 and MOP3, and index differences from the content frames of those newly created MOPs. The meeting with Dayan and the meeting about Jerusalem are indexed in MOP1, and also in appropriate specialized E-MOPs.

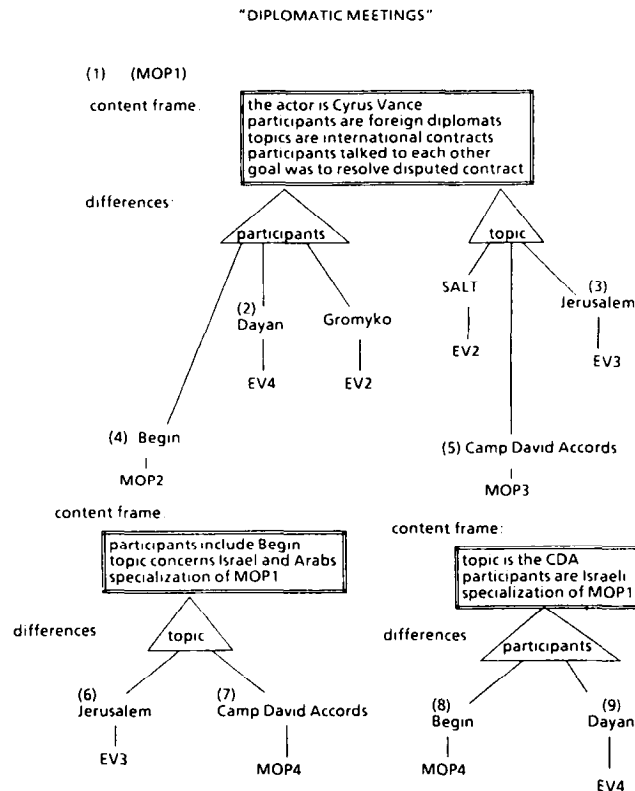


Figure 6

3. TRAVERSING AN E-MOP— THE BASIC RETRIEVAL PROCESS

Because events are organized in conceptual categories by their differentiating features, an event can be found in an E-MOP by following indices corresponding to its features. This process is a *traversal* process. Traversal involves following appropriate indices down the tree until the sought item is found. An event to be retrieved from an E-MOP is called a *target* event, and the features which describe it make up its *context specification*. A target event, or event which must be retrieved, can be said to be *targeted* for retrieval.

Traversal of an E-MOP is guided by the particular event targeted for retrieval. Consider, for example the following questions:

- (Q2): Have you ever discussed SALT with Gromyko at a diplomatic meeting?
- (Q3): Have you ever attended a diplomatic meeting about the Camp David Accords with Dayan?

As a first step, answering a question requires extraction of its “target concept,” that is, the concept that must be searched for in memory. The target concept for (Q2) is a “diplomatic meeting with SALT with Gromyko.” Answering (Q2) requires retrieval of that event.

The next step in the process involves choice of indices. Index selection is based on features specified in the target event. Indices chosen for use in retrieving any target event should be features that would have been chosen as indices for that event if it had previously been indexed in the E-MOP, i.e., features which would have differentiated it from other events already in the E-MOP. (The actual process of choosing indices will not be discussed here, but can be found in Kolodner, 1981a, 1983, 1984.)

A “diplomatic meeting about SALT with Gromyko” can be retrieved from the structure in Fig. 6 by traversing either of the indices “has topic SALT” or “has Gromyko as a participant,” retrieving the event found at each of those points (EV2) and checking to make sure it has all the required features. Since EV2, found at both index points, is a meeting with Gromyko about SALT, it has all the features of the target event and can be used to form an appropriate response to the question.

When a target event specifies an event feature which is unique in an E-MOP, the event can be found by traversing the index associated with that feature. Question (Q2) had two unique features in the “diplomatic meetings” MOP—its participants and its topic. Either can be followed to retrieve the appropriate event. The target concept for (Q3) is a “diplomatic

meeting with Dayan about the Camp David Accords.” It can be retrieved from Fig. 6 by traversing the index corresponding to the unique MOP feature “has Dayan as a participant” (2). EV4 would be found. It would be checked to make sure it had the topic “the Camp David Accords” (the remainder of the target concept’s specification). Since EV4 is a meeting with Dayan about the Camp David Accords, it has all specified features and can be used to answer the question.

There is also a second way the “meeting with Dayan about the Camp David Accords” can be found in that tree. One feature that meeting has is “has topic the Camp David Accords.” If that index were traversed, MOP3, “diplomatic meetings about the Camp David Accords,” would be reached. When a specialized E-MOP is reached during the traversal process, it must be searched for the target event in the same way its parent E-MOP is searched. Thus, its indices are traversed after selecting appropriate indices for traversal to find the target event. In this case the index “has Dayan as a participant” (9) of MOP3 would be traversed and EV4 would be found, again a sufficient answer to the question.

Because there is no way of knowing before traversal whether or not a feature is unique to an E-MOP, the index associated with each feature selected for retrieval must be traversed. If one is unique, an event will be found and traversal can end. Otherwise traversal continues at the next E-MOP level. Thus, in answering (Q3), both indices “has Dayan as a participant” and “has topic the Camp David Accords” are traversed. Since one is unique, an event is found, and traversal can stop. If, however, the index “has Dayan as a participant” had not been unique (i.e., if there had been more than one meeting with Dayan indexed in the E-MOP), then traversal would have had to continue with the E-MOP at that point and within the “meetings about the Camp David Accords” MOP.

Thus, traversal is a *recursive process* involving choice of indices and traversal of those indices. It stops when an event is found or when there are no more specified indices to be traversed. If there are multiple paths to a target event, it will be retrieved from the shortest path that has all of its indices specified in the target event. We can think of traversal as a breadth-first search which implements parallel traversal of all appropriate indices. Figure 7 specifies the algorithm for traversal.

4. STRATEGIES FOR SEARCHING MEMORY

If we assume that traversal is the basic or core retrieval process, then the entire retrieval process can be defined as follows: Choose a conceptual category for search, and traverse the indices of that category to find the target concept. Three kinds of failures derive from this algorithm:

E-MOP Traversal

1. Select possible indices for the target event based on its specified features and their differences from the norm of the E-MOP being traversed.
2. IF there are no indices, THEN return "not found".
3. ELSE follow all of those indices in the E-MOP.
4. IF events are found, THEN check that they have all features of the target event. If any do, return them, and finish.
5. IF E-MOPs are found, traverse them in parallel using this algorithm.

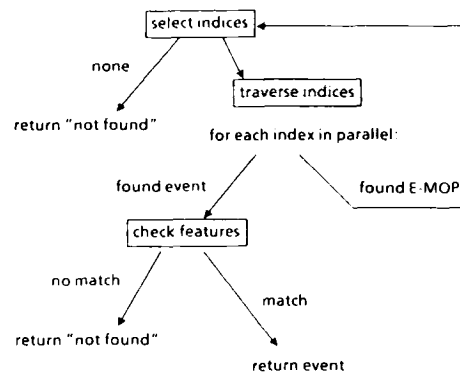


Figure 7

1. A question may not specify a category for search. In that case, one must be chosen.
2. Features specified in a question may not match those indexed in memory. In that case, hypothetical features must be derived.
3. A target event (extracted from the question) may describe many rather than only one event in memory. Because enumeration is not allowed, there must be some other way of getting at individual events. If the process in 2 is unsuccessful, then hypothetical related events must be derived and searched.

Our theory of reconstructive memory will present retrieval strategies to correct each of these failures. Figure 8 shows where each fits into the traversal algorithm described, and shows the entire retrieval process. Boxes corresponding to the three cases above are marked. In the remaining sections, the components of this process and their implications will be described.

4.1 Choosing a Category for Search

In order to search a memory organized around conceptual categories according to the traversal algorithm given in Fig. 7, the category or categories

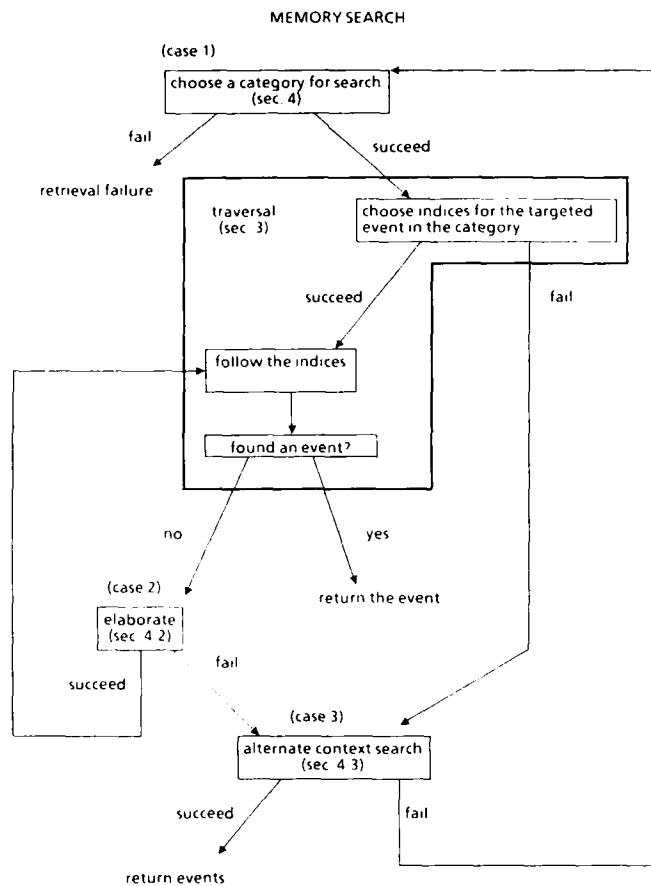


Figure 8

to be searched must first be specified. Consider, for example, the following question:

(Q4): Mr. Vance, when was the last time you saw an oil field in the Middle East?

If "seeing oil fields" were one of memory's categories or E-MOPs, then this question would be fairly easy to answer. "Seeing oil fields" would be selected for search. If it indexed an episode in the Middle East, that episode could be retrieved from it. Similarly, if "seeing objects" were a memory category, it could be selected as a category for retrieval, and events in the Middle East and at oil fields could be retrieved from it.

Suppose, however, that neither of these categories existed in the memory being searched. In that case, an E-MOP for search would have to be

chosen. We can imagine that the following reasoning process might take place during the process of answering (Q4):

A1: An oil field is a large sight; perhaps I saw an oil field during a *sightseeing* episode in the Middle East.

If one used information about episodic contexts associated with “large sites” a “sightseeing” category could be chosen for retrieval. Its contents could be searched for an episode at oil fields in the Middle East. If episodes in the sightseeing category were organized by the type of site and part of the world where they took place, and if there had been an episode in the Middle East at an oil field, then “a sightseeing episode at an oil field in the Middle East” could be retrieved.

The need to choose a category for search is common to both artificially constructed databases and human memory. In a complex database, such as one which stores many different kinds of events, we should not expect a user to know all of memory’s categories. Nor can we expect that every natural language query asked of a database will specify a category or record type for search. The retrieval processes must therefore be prepared to generate categories for retrieval. To see that this process also occurs in human memory access, consider this question and a typical response:

(Q5): Who is the most famous person you have ever met?

The following protocol is typical of the answers to this question:

First, I thought how somebody could be famous, and politics was the first thing I thought of. Then I thought about circumstances where I could have met a famous politician. I searched for political experiences I have had—mostly political rallies I participated in and experiences campaigning for candidates. I remembered that I had met McGovern. But since you said “most famous,” I went on to think of other famous people I might have met. Next I thought of entertainers, and how I could have met them. I remembered going backstage a few times after seeing shows, but I couldn’t remember whom I had met. I started going through TV programs and the entertainers on those programs. I couldn’t think of meeting any of those. Then I thought of famous scientists, and where I could have met them.

The diverse kinds of situations recalled in answering this question suggest that this person had no conceptual category (E-MOP) for “meeting famous people.” Rather, he seemed to use knowledge about “famous people” to choose types of situations in which he could have met one. Using that information, he was able to direct search for episodes where he might have met politicians, entertainers, and scientists. His knowledge about politicians directed his search toward political rallies and campaigns where he

could have met politicians. His knowledge about actors directed the search toward times he had been backstage after a show.

To find out how search can be directed toward the right kinds of situations, consider the inferential knowledge which must be present in memory to direct search in the protocol quoted:

1. Politicians take part in political rallies.
2. Politicians take part in political campaigns.
3. Actors perform in theaters.
4. After shows, actors are backstage at theaters.
5. Actors perform in TV shows.

Each of these statements associates a class of people (e.g., politicians, actors) with a type of situation (e.g., political rallies, going to the theater). If each item in memory describes the event contexts (i.e., conceptual event categories) in which it normally occurs, then appropriate categories for search can be chosen by selecting from the concepts associated with event components specified in or inferred from a question.

Memory Principle #5

Retrieval from memory requires knowledge about the contexts associated with target items.

If this is so, then memory search can be constrained and directed only to relevant categories. A conceptual category chosen for retrieval provides a context for search. "Political rallies" can be chosen as a context for search because they are associated with "politicians," one kind of famous person. "Sightseeing" can be chosen as a context for "seeing oil fields," if "oil fields" are defined as a "large site" and if "large sites" are associated with "sightseeing." The strategies for choosing contexts for retrieval are called *component-to-context instantiation strategies*. To see how these strategies work, consider the following examples:

- (Q6): Have you ever discussed the Camp David Accords with Dayan?
 (Q7): Has your wife ever met Mrs. Begin?

These questions are appropriate to ask the CYRUS data base about Vance. Suppose we wanted the CYRUS data base to answer these questions. CYRUS knows about many different situations in which Vance might talk to other people (in fact, almost any situation could apply), but it does not have an E-MOP for "talking to people." Although Vance could talk to somebody during any of his activities, it would not be appropriate for CYRUS to

search every one of its E-MOPs to find discussions about the Camp David Accords or between his wife and Mrs. Begin.

In general, in order to search memory, a *context for search* must first be set up. A context for search specifies a memory category to be searched. Because CYRUS's memory is organized around event categories, or E-MOPs, a context for search in CYRUS must include a specification of an E-MOP for traversal. One way people's memories are organized is around events. Thus, in choosing a context for search, a person might also choose event categories. People's memories are also organized in other ways, and a person might choose any of those types of categories in this step.

The point is, if a question does not already mention an available context for search, then one must be derived. That is what these strategies do. They are applied as the first step of memory search if a retrieval request (e.g., question) does not already refer to a category which can be traversed. CYRUS chooses appropriate E-MOPs for traversal by looking at the kinds of situations associated with each of the given question components. Thus, these strategies are chosen by examining question components and choosing strategies relevant to the specified components.

In answering (Q6), the topic and participants, for example, are used to infer a context for search. "Infer-from-Topic" makes use of topic information to infer an E-MOP, while "Infer-from-Participants" makes use of participant information. These two strategies are listed in Figs. 9 and 10.

INFER-FROM-PARTICIPANTS

If particular participants, or classes of participants, such as people, groups, or organizations, are specified, and if any has a particular event context associated, predict those characteristic events associated with each that could also include other already-specified components.

Figure 9

INFER-FROM-TOPIC

If topics of discussion or classes of topics of discussion are specified, and if any of them have communicatory event contexts associated with them, predict appropriate communicatory events associated with the topic or class of topic.

Figure 10

Since context-instantiation strategies infer event contexts from event components, memory must have information associating each of its contracts and people with E-MOPs. In particular, in order to use the topic and persons in Q6 to derive a context for search the "Camp David Accords" is identified as an "international contract," and that in turn, has the context

“political meeting” associated with it. Similarly, Dayan is identified as a “foreign diplomat,” which has the event context “diplomatic meeting” associated with it. Figure 11 shows CYRUS’s organization of information about international contracts.

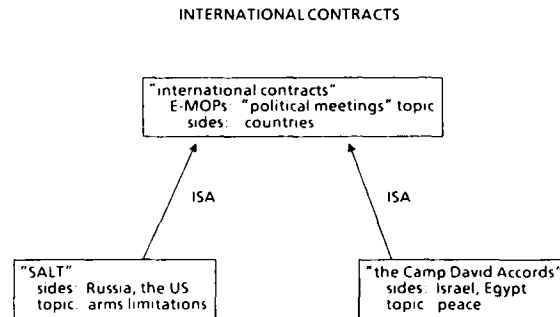


Figure 11

“International contracts” lists “political meetings” as an E-MOP it is associated with, and further states that it fits the topic slot of such a meeting. Using this information, “Infer-from-Topic” chooses the context of a “political meeting” to answer Q6. Using similar information associated with “foreign diplomat” (of which Dayan was one, “Infer-from-Participants” chooses “diplomatic meeting,” a refinement of “political meeting” as the context for search. After applying these strategies, CYRUS searches memory for a “diplomatic meeting about the Camp David Accords with Dayan.” The following example shows CYRUS using “Infer-from-Topic” to choose a context for answering a question similar to (Q6).

```

Enter next question:
>Who have you discussed the Camp David Accords with?

The question is
((ACTOR HUM1 < = > (*MTRANS*) MOBJECT
  (*CONCEPTS* CONCERNING CNTRCT1) TO (**)) TIME G1380)
The question type is "concept completion"

The question concept is
((ACTOR HUM1 < = > (*MTRANS*) MOBJECT
  (*CONCEPTS* CONCERNING CNTRCT1) TO (**)) TIME G1380)

  applying INFER-FROM-TOPIC to CNTRCT1
  inferring undifferentiated political meeting

The inferred question concept is
((( < = > ($M-MEETING ACTOR HUM1 OTHERS (**)) TOPIC CNTRCT1))
  TIME G1380)

searching memory for question concept
  searching directly for $M-MEETING
  searching for $MEET
  searching for $CONSULT
  searching for $PUB-REL-MEET
  ...

Begin, Sadat, Dayan, Hussein, and Assad at diplomatic meetings, Carter during
consultations, and a group of Jewish leaders at a public relations meeting
  
```

Figure 12

The strategy “Infer-from-Participants” is used to answer Q7. CYRUS’ use of this strategy is illustrated in Fig. 13. In this case, CYRUS uses knowledge about the relationship between his wife and Mrs. Begin, and himself and his wife. CYRUS determines that if it knew of Vance’s wife and Mrs. Begin’s meeting, it must have been a political social event in which the two husbands were also involved. It finds its answer by searching for “political social events” it knows about and finds a “state dinner” in Israel.

```

Enter next question
> Has your wife ever met Mrs. Begin?

The question is
((ACTOR HUM4 IS (*PROX* VAL G1571)) TIME G1572 MODE (**))
The question type is "verification"

The question concept is
((ACTOR HUM4 IS (*PROX* VAL G1571)) TIME G1572)

  applying INFER-FROM-PARTICIPANTS
    to family group (HUM1 HUM4)
    inferring social occasion
  applying INFER-FROM-PARTICIPANTS
    to occupation group (HUM1 HUM60)
    inferring political occasion
    inferring political social occasion

The inferred question concept is
((<=> (SM-SOCIAL-OBL ACTOR HUM1 OTHERS (HUM4 HUM60 G1571)))
 TIME G1572)
searching memory for question concept
  searching directly SM-SOCIAL-OBL
  searching for $PARTY
  searching for $DINE
  searching for $STATE-DINNER
...
yes, most recently at a state dinner in Israel

```

Figure 13

In CYRUS, the types of concepts which have E-MOPs associated with them are those concepts (and classes of those concepts) which are associated with events. In CYRUS’s memory, that includes animate participants, such as people, organizations, countries; inanimate participants, such as objects; places; times; and topics. In addition, the classifications each particular person, country, or object fits into (e.g., “foreign diplomat”) have E-MOPs associated with them. In general, automatic context construction in any data base requires that data base concepts refer to the types of valid retrieval contexts they can be associated with.

4.2 Choosing Indices During Retrieval

Traversing a category’s indices to find a target event requires selecting potential indices from the event’s features and traversing those indices. Features specified in a question, however, may not correspond to those indexed in memory. Just as we cannot expect a questioner to know all of the categories represented in a memory, neither can we expect him to know another memory’s indexing scheme, whether it is a computer memory or a natural

human memory. We must therefore look for some means of automatically deriving features which are indexed in memory from others a questioner might specify.

Consider again, the question earlier cited:

(Q8): Mr. Vance, when was the last time you saw an oil field in the Middle East?

We mentioned earlier that if a “sightseeing” E-MOP organized its episodes according to their geographical location or the type of site, then finding a sightseeing event in the Middle East would be easy. Suppose, however, that our “sightseeing” category does not organize its episodes by either of those features. In that case, the following *elaboration* of the event sought in the question might be appropriate to answer the question:

A2: Which countries in the Middle East have oil fields? Iran and Iraq have oil fields and so does Saudi Arabia. Have I seen an oil field in any of those places?⁶

If sightseeing episodes are organized by country, then *elaborating* on “the Middle East” and specifying particular countries in that area of the world would enable retrieval of sightseeing episodes that took place in each of those places. Instead of searching for “sightseeing at an oil field in the Middle East,” one could search for each of the more specific episodes, i.e., “sightseeing at an oil field in Iran” or “sightseeing at an oil field in Iraq.” This process which transforms given features of a target event into features which might be indexed in the current E-MOP is called *index fitting*. Since the values of indices are not enumerable, this process is actually one of elaborating plausible features of sought items in the hope that one will be an index and “unlock” the network. Index fitting allows *downward traversal* in memory’s networks to find nodes representing more specific concepts.

Elaboration is also necessary when a target concept is *too general*, and no additional features are available to use for traversal. In that case, we call the elaboration process *plausible-index generation*. Consider, for example, the following human protocol of a person recalling the museums he had visited:

I know I’ve been to a lot of museums in Europe. I’ve been to England, and I went to a number of museums there—some in London—the British Museum, the National Gallery, and a few smaller galleries. I was at a

⁶Note that while we have prohibited enumeration of events from memory categories, this answer requires enumeration of countries in the Middle East. We do not address the process of such enumeration here. The way this could be done is by imagining a map of the Middle East and mentally walking around that map. Alternatively, it could be done by reconstructing this map. Thus, blind enumeration (i.e., enumeration of a list) is not necessary.

museum in Brighton—the Royal Pavilion. I've been to museums in Paris—the Louvre and some smaller ones. In Rome. I've been to In Naples, to In Florence, to

Previously, we observed that this person was trying to better specify or elaborate his descriptions of museum visits so that he could remember them one at a time. Because the person had been to many museums, however, "museum visit" is too general a specification of a target event—it describes too many events. The traversal process described earlier must fail, because no features are available for it to use as traversal indices. In this case, elaboration provides additional hypothetical features of the target event. Thus, we can explain that this person was generating additional cues for downward traversal of memory structures.

Index fitting and plausible-index generation are used to elaborate features of a target event with respect to a particular E-MOP. Those E-MOPs are the ones at which traversal cannot continue. Thus, different plausible elaborations for a particular event might result when it is considered in two different MOP contexts.

The retrieval strategies used for elaboration are called *instantiation strategies*. They are applied under the following conditions:

1. Traversal has failed because there are no additional features to use as indices (i.e., the target event is too generally specified).
2. Traversal has failed because the only features available for use as indices do not correspond in type to E-MOP indices.

The first case results in plausible-index generation. The second results in index fitting. Note that if available features correspond to E-MOP indices in type but not value, then further work need not be done—we know the target event is not in memory.

The museum example above illustrates the need for elaboration in human memory. To see how it fits into our own model consider the following question asked of the MOP in Fig. 6:

(Q9): How many times have you met with Begin about the Camp David Accords?

One way to answer this question is to find all meetings with Begin about the Camp David Accords. Because enumeration of events is impossible, answering this question requires individual retrieval of each appropriate meeting. In order to traverse E-MOPs to find those meetings, the indices to be traversed must be specified.

Using the features already specified in the question, traversal can proceed as far as MOP4, the sub-MOP which organizes "diplomatic meetings with Begin about the Camp David Accords." Because there are no addi-

tional features specified in the target concept which correspond to indices in that E-MOP, however, the traversal process must abort at this point. Although all meetings between Vance and Begin about the Camp David Accords can be found in MOP4, they can only be retrieved through specification of the differences that index them. The retrieval specification “meetings with Begin about the Camp David Accords” is too general and plausible-index generation must be done.

Suppose that MOP4—“diplomatic meetings about the Camp David Accords with Begin”—has the structure shown in Fig. 14. If possible places for these meetings or the types of events they were part of could be inferred, then actual meetings could be retrieved from the MOP.

‘meetings with Begin about the Camp David Accords’ --MOP4

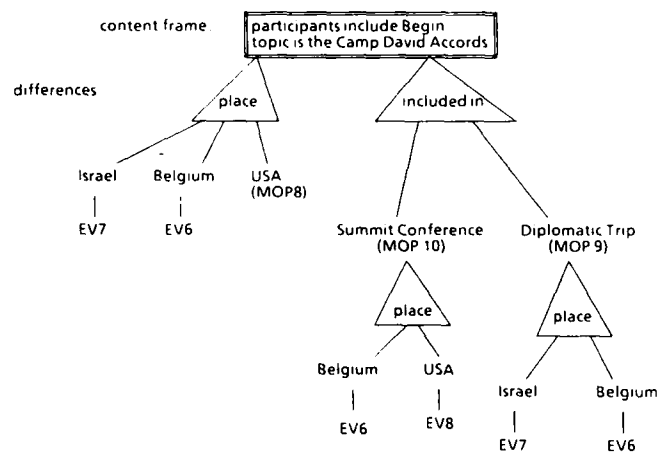


Figure 14

Instantiation strategies are used to elaborate on two classes of event features: event components and event contexts. Event components include such descriptive features as participants, location, and topic. An event's context includes other episodes related to it through time, causality, or containment. Two classes of instantiation strategies are used to do this elaboration: *component-instantiation strategies* elaborate on event components, while *context-to-context instantiation strategies* are used to infer an event's context. Each will be described in turn.

4.2.1. Component-instantiation strategies. Consider the instantiation rule in Fig. 15. Applying this rule to “meetings with Begin about the Camp David Accords” (the target event of (O9)), “Israel” can be inferred as a plausible place for one of these meetings. The index for “location is Israel”

INFER-COUNTRY

To infer the country an event might have taken place in, use the participants' country or residence, country they habitually travel to, or their nationality.

Figure 15

can then be traversed in the E-MOP, and EV7 will be found. The set of strategies which infer event components and slot fillers are called *component instantiation strategies*.

These strategies use two types of knowledge to infer plausible event components:

1. content-frame information describing default values for components of items in that E-MOP.
2. the general relationships (less dependent on content) between components and other components.

For an example of the first type of knowledge, see the role filler information associated with "diplomatic meetings" in Fig. 3. One specification associated with "diplomatic meetings" is that its participants are foreign dignitaries of countries involved in the contract being discussed. Such MOP-specific information is used to constrain application of more general relationships between components. Using this information, the participants in a meeting about the Camp David Accords can be inferred to be Israeli or Egyptian diplomats. Similarly, the knowledge that a "diplomatic meeting's" location is a conference room in the capital city of the country of residence of some important participant can be used to infer the location of a "meeting with Begin." It could plausibly have been in Israel, and more specifically Jerusalem.

General relationships between components are taken into account after MOP-specific constraints are applied. If a "ski trip" MOP, for example, specified that its location was mountains, then a component-instantiation strategy for enumerating mountains (perhaps by walking around a map) would be applied after the MOP-specific constraint was made. Note that both types of rules access specific details of the target event.

In memory, component-instantiation strategies are associated with the types of components they elaborate on. Strategies for participants are stored in a node in memory describing "participants." The strategy for inferring the country an event may have occurred in is associated with a memory node for "country" which is accessed through a more general strategies associated with "location." That strategy calls "Infer-Country" when the location to be inferred is already constrained to be a country.

4.2.2 Context-to-context instantiation strategies. In addition to “location” indices, the “meetings with Begin about the Camp David Accords” MOP (specified in Fig. 14) has indices for types of E-MOPs its events were included in. Thus, elaborating the kinds of episodes a “meeting with Begin about the Camp David Accords” was part of can also help in retrieving individual events from that E-MOP. The set of instantiation strategies which elaborate on contexts related to a target event are called *context-to-context instantiation strategies*. These strategies use content frame information from E-MOPs the target event fits into to produce descriptions of events plausibly related to the target event.

Context-to-context instantiation strategies produce a set of event descriptions as output. There is one such strategy corresponding to each possible relationship between events. The set of context-to-context instantiation strategies is listed in Fig. 16. Input to these strategies consists of an event (real or plausible) and an E-MOP that event fits into. The event descriptions output by the strategy describe events plausibly related in the specified way to the input event.

CONTEXT-INSTANTIATION STRATEGIES

Instantiate-Enablements
 Instantiate-Preconditions
 Instantiate-Results
 Instantiate-Reasons
 Instantiate-Enabled-Events
 Instantiate-Larger-Episodes
 Instantiate-Seq-of-Events
 Instantiate-Preceding-Events
 Instantiate-Following-Events
 Instantiate-Standardizations

Figure 16

Consider, for example, the strategy “Instantiate Larger Episodes” (see Fig. 17). It produces descriptions of episodes an event might have been part of. Each strategy has three steps. First the appropriate content frame component of E-MOP is retrieved, e.g., larger episodes for “Instantiate-Larger-Episodes,” resulting episodes for “Instantiate-Results.” Second, role fillers or components of EVENT are transferred appropriately to the event framework provided by execution of the first step. Third, component-instantiation strategies are applied to further refine the role fillers of EPISODES.

Applying this strategy to a “meeting with Begin about the Camp David Accords” (the target event of Q9), and using content frame knowledge from the “diplomatic meetings” E-MOP (see Fig. 3), the following descriptions are generated of episodes the meeting could have been part of:

INSTANTIATE LARGER EPISODES

EVENT = the event whose larger episodes are to be instantiated

E-MOP = the E-MOP EVENT may fit into

EPISODES = the newly-instantiated events (i.e., the episodes which may plausibly contain EVENT)

- (1) Get the "larger episodes" of E-MOP.
- (2) For each larger episode specification, transfer components as specified, producing EPISODES, and make the following inferences when more specific information is not available on the E-MOP:
 - (a) Time specifications on each EPISODE include the time specified in EVENT.
 - (b) Place specifications on each EPISODE include the place specified in EVENT.
 - (c) Participants of EVENT are included in the participants of each EPISODE, groups and organizations they belong to might also be involved.
 - (d) Topic of each EPISODE includes that of EVENT.
- (3) Use relevant Component-Instantiation strategies to further specify components of each EPISODE, using the already-instantiated partial descriptions of each component as constraints.

Figure 17

- A "diplomatic trip" to Israel.
- A "Summit Conference" with Begin as a participant and whose topic was related to the Camp David Accords.
- "Negotiations" with Israel about the Camp David Accords.

Production of these episode descriptions allows the following elaborations of that target meeting:

- "meetings with Begin about the Camp David Accords which took place during a diplomatic trip to Israel"
- "meetings with Begin about the Camp David Accords during a summit conference which included Begin as a participant and whose topic was related to the Camp David Accords."
- "meetings with Begin about the Camp David Accords which were part of negotiations with Israel about the Camp David Accords"

Retrieval of each of these can then be attempted. The first will retrieve EV7 by following the indices "included in a diplomatic trip" and then "place is Israel" in the E-MOP is Fig. 6. Retrieval of the second will include

traversal of the index “included in a summit conference,” and further elaboration will be needed to retrieve an individual event. The third will not be successful in the E-MOP illustrated.

Figure 18 shows an example of CYRUS using instantiation strategies.

```

Enter next question:
>Who have you discussed the Camp David Accords with?

The question is:
((ACTOR HUM1 < = > (*MTRANS*) MOBJECTS
  (*CONCEPTS* CONCERNING CNTRCT1 TO (*?*))
The question type is "concept completion"
The question concept is:
((ACTOR HUM1 < = > (*MTRANS*) MOBJECTS
  (*CONCEPTS* CONCERNING CNTRCT1 TO (*?*))
  applying INFER-FROM-TOPIC to CNTRCT1
  inferring political meeting context
The inferred question concept is:
((< = > (sM-MEETING ACTOR HUM1 OTHERS (*?) TOPIC CNTRCT1)))
searching for question concept
  searching directly for sM-MEETING
  additional information needed

applying strategies to elaborate question concept
...
elaborating others
  OTHERS are RT-DIPLOMAT of Israel, Egypt, or the USA
  searching for sM-MEETING with OTHERS = RT-DIPLOMAT of Israel
  additional information needed
  elaborating OTHERS
    OTHERS could be Begin, Dayan, or Weizmann
    searching for sM-MEETING with OTHERS = Begin
  ...
  searching for sM-MEETING with OTHERS = Dayan
  ...
  searching for sM-MEETING with OTHERS = Weizmann
  ...
  searching for SM-MEETING with OTHERS = RT-DIPLOMAT of Egypt
  additional information needed
  elaborating OTHERS
    OTHERS could be Sadat
    searchng for SM-MEETING with OTHERS = Sadat
  ...
  ...

Israeli and Egyptian diplomats including Begin, Dayan, Weizmann, Sadat, and
Khalil, and American diplomats including Carter and Brown.

```

Figure 18

4.3 Searching for Something Other Than What Was Asked For

One of the observations we can make from many of the examples and protocols above is the following:

Memory Principle #6

Retrieval often requires search for something other than what was requested.

Two reasons why it is often necessary to search for something other than what was requested are that (a) a category for search might not be specified in a question, and (b) a given specification might be insufficient for retrieval. In the second case, the given context is elaborated as described above, and a better specified context is searched for. Elaboration of plausible features is not always successful, however. Unsuccessful elaboration is a third reason for searching for something other than what was requested.

Consider, again, our sightseeing example from the previous section:

(Q10): Mr. Vance, when was the last time you saw an oil field in the Middle East?

Suppose that there were not enough information available to choose indices for traversal. In a memory where contextually related records refer to each other, another way to find a sought item is to search for items it might have been related to. Another related item might be more easily retrieved, and if found, might refer to the targeted one. Since any sightseeing episode in the Middle East would have had to have happen during a trip to the Middle East, the following reasoning would be appropriate to recall "sightseeing episodes in the Middle East."

A3: In order to go sightseeing in the Middle East, I would have had to have been on a trip there. On a vacation trip. I wouldn't go to see oil fields, so I must have been taken to oil fields during a diplomatic trip to the Middle East. Which countries might have taken me to see their oil fields? Saudi Arabia has the largest fields, perhaps they took me to see them. Yes, they did when I was there last year.

Search of this kind is quite often done by people, as shown in the following example of a person naming the museums he had been to:

Let me see. What other museums have I been to? The last time I was in England, I went to a bunch of museums in London—the British Museum, some gallery whose name I can't remember, and Mme. Taussaud's Wax Museum. And I also went to some palaces there that were museums—the Royal Pavilion in Brighton and some other palace in London, um, —it was called Hampton Court. When I went on my first trip to France, . . . I once went on a trip around New York State, and I went to the photography museum in Rochester and the Corning Glass Museum. During my trip to California . . .

In each of these examples, we see search for a different kind of event than the one originally asked for. To find sightseeing experiences in A3, we see search for diplomatic trips during which the sightseeing might have taken

place. In the human protocol, we see the person recalling trips to remember museum visits. Searching for a different type of event than the one targeted for retrieval is called *alternate context search* (case 3 in Fig. 8). Once a related event is retrieved, it can provide cues to use in reconstructing the originally targeted event. Features of a trip to England, for example, can be useful in filling in details of museum visits experienced during the trip.

While index fitting and elaboration direct depthwise search within memory, alternate context search is *lateral traversal* of memory. Search is within different E-MOPs, or contexts, than the one originally selected for retrieval. One problem associated with searching for alternate contexts is constraining the search. Search should be constrained to contexts that have a possibility of being related to an item targeted for retrieval. In general, for search to be constrained to relevant contexts, memory categories must hold *generalized information* concerning the relationships of their items to items in other memory categories. If “museum visit” holds the information that its items are often parts of “trips” or “visits from out of town friends” or “visits to New York”, then search can be constrained to only those three alternate contexts. Similarly, for constraining search of contexts related to “diplomatic meetings,” generalized information associated with diplomatic meetings must be used.

*4.3.1. Guiding alternate-context search. Executive strategies*⁷ direct search for alternate contexts. They have three steps: (1) they first call context-to-context instantiation strategies (explained earlier) to construct descriptions of episodes plausibly related to the target event; (2) they then pass control to the traversal/elaboration process, directing that process to attempt retrieval of each of those hypothetical events; and (3) if any such event is found, its episodic context is searched for the original target event.

To search for a particular “diplomatic meeting,” for example, executive strategies trigger appropriate context instantiation strategies to construct contexts related to diplomatic meetings (e.g., negotiations, summit conferences, crisis situations, diplomatic trips), traverse memory to retrieve those hypothetical related contexts, and if one is found, direct search within the episodic context of that episode to find the targeted diplomatic meeting.

Steps 1 and 2 of executive strategies have been described previously. Search of surrounding context (step 3) must be described in more detail. Recall from the discussion of instantiation strategies that an event’s context includes events it was related to through containment, time, or causality. Search of the episodic or surrounding context of one event for a target event means reconstructing the appropriate part of the first event to see if it corresponds to the target event. In our model there are two mechanisms for doing this. First, an event itself might point to or describe another event related to it. If a particular museum experience was the highlight of some

⁷In previous papers, these strategies were called search strategies.

trip, for example, then the trip's description in memory may point directly to the museum experience or describe it in enough detail to retrieve it from its own MOP. In this case, to find a target event in the surrounding context of another event, the possible relationships between the two events is determined, and appropriate parts of the context of the second event are checked to see if events it points to or describes correspond to the target. If an event does not explicitly point to episodes related to it, a second procedure is available. The content frame of the E-MOP that event resides in might specify the types of events related to it. "Diplomatic meetings," for example are often parts of "conferences." In that case, the content frame description of a hypothetically related event is used by context-to-context instantiation strategies to reconstruct that aspect of the event's context. Again, correspondence to the target event is checked.

"Find-from-Larger-Episodes" (see Fig. 19), for example, constructs, finds, and searches the episodic context of events a target event could have been part of. To aid in finding a particular "diplomatic meeting with Gromyko," it would (1) call the context instantiation strategy "Infer-Larger-Episodes" to construct contexts for "summit conferences about a Russian-American concern," "diplomatic trips to Russia," and "negotiations concerning a Russian-American contract," (2) traverse memory searching for each of those, and (3) search (using the methods described in the previous paragraph) the sequence of events of each episode found for an appropriate diplomatic meeting.

FIND-FROM-LARGER-EPIISODES

IF the Target Concept could have been embedded in a larger episode

THEN

- (1) Use "Instantiate-Larger-Episodes" to hypothesize episodes the target concept could have been part of. Use Component-Instantiation strategies to fill in each description.
- (2) Search memory for those instances
- (3) If one is found, search its sequence of events and instrumental events for events which could match the Target Concept.
- (4) If no match is found, use "Instantiate-Seq-of-Events", plus the additional information found in the larger episodes, to better specify the target concept. Traverse memory for the new target concept.

Figure 19

If an event plausibly related to the target event is found, but it does not point to an event corresponding to the target event, it can nevertheless aid in retrieval of the target. This second way of using alternative events to find a target is a process of using features of the alternative event to flesh out the target. Its use in "Find-from-Larger-Episodes" is listed in step 4 of that strategy. As an example of its use, consider the following: The target

event is “museum experiences.” In searching for museum experiences, some trips are recalled, for example, a trip to England. While the sequence of events of the trip itself may not refer to any museum visits (step 3), transference of its features to a “museum visit” context might aid in describing a museum visit that did in fact happen. Its time, place, the weather, etc., applied to a museum visit might provide enough details to retrieve a valid visit to a museum.

In guiding search for event contexts related to a target event, search strategies direct search for “what must have happened” or “what may have happened” if the target event had taken place. If an event that must have happened along with a targeted event can be found, its specification in memory might refer to the target or aid in its reconstruction. There is an executive search strategy corresponding to each relationship an event can have to other events. The entire list is in Figure 20.

EXECUTIVE SEARCH STRATEGIES

Find-from-Enablers
 Find-from-Preconditions
 Find-from-Results
 Find-from-Reasons
 Find-from-Enabled-Events
 Find-from-Larger-Episodes
 Find-from-Seq-of-Events
 Find-from-Preceding-Events
 Find-from-Following-Events
 Find-from-Standardizations

Figure 20

4.3.2 When is alternate-context search useful? Alternate-context search is helpful under two circumstances: (a) when the event being searched for is obscure, and (b) when many events corresponding to a general specification must be found.

To see what is meant by an “obscure event,” consider the following question:

(Q11): Have you ever rented a Chevy?

Suppose a person who had rented cars many times were answering this question. People renting cars normally rent by size and not by manufacturer. Thus, unless a person once had a particularly bad or good experience renting a Chevrolet, he would probably not be able to retrieve an experience renting one without further specification. When he attempted retrieval, he might recall the following: “I’ve rented cars many times, and I need more information to find recall renting a Chevy.” If he cannot elaborate on additional features of “Chevy renting,” then he will have to answer “I don’t

know," although the event might be in memory. This is an example of an obscure event. It is one which might be present in memory, but for which there is not enough generalized information available to do the necessary elaboration to get it out.

When retrieval fails because of inability to produce a sufficient elaboration, executive strategies are applied to search for alternate contexts related to the event which might be more easily retrieved. In one case, we observed a person doing the following reasoning in order to recall if she had ever rented a Chevy:

I rent cars during trips when I need to be able to travel easily. In California, it is impossible to get around without a car. On my trip there last summer, I rented a Fairmont, and on my last trip there, I had a Toyota. Neither of them was a Chevy. What kinds of bad experiences have I had with rental cars? Was one of those cars a Chevy? I once missed a plane because I couldn't find the rental car place to return my car—that was a foreign car of some kind. Another time, I rented a car that was too big for me—I couldn't see over the steering wheel. It was the only car they had available. That might have been a Chevy. I have no idea if it was or not.

In giving this answer, the respondent recalled trips during which she might have rented cars and rental car experiences during those trips. She then recalled bad experiences while renting a car to see if any of those involved Chevrolets. In other words, she searched memory for episodes related to renting cars.

Searching for alternate contexts can aid retrieval since an alternate context might be less common than a target event, and thus need less specification in order to be retrieved. Summit conferences happen less often than diplomatic meetings do. Although a specification "diplomatic meetings about SALT" might describe many meetings, the specification "summit conferences about SALT," which can be derived from that, describes only a small number of events.

Alternate-context search can also aid retrieval of a target event if the alternate context can be better specified with respect to the E-MOPs it is indexed in than the target event can. "Diplomatic meetings with Begin" might be an ambiguous target event. Although "diplomatic meetings" might often happen as a result of "crisis situations," if there had only been one crisis situation which involved Israel, then it would be better specified in its E-MOPs than "meetings with Begin" are in "diplomatic meetings." Thus, it would be more retrievable and aid retrieval of at least one meeting with Begin.

In extensive memory search, executive strategies can be applied in succession until a satisfactory answer is retrieved from memory. When people attempted to recall museum experiences, they seemed to apply search strate-

gies in succession. In their observations of people recalling persons in their high school class, Williams and Hollan (1981) reported application of a succession of search strategies until a sufficient answer was found. In attempting to recall meetings with Gromyko, CYRUS recalls trips to Russia, negotiation episodes involving SALT, and summit conferences about SALT and arms limitations. Memory search continues until strategies stop producing results or until a sufficient answer is found.

When CYRUS searches its memory extensively, it applies each appropriate strategy in succession until it has found a satisfactory answer. Thus, when searching for all of Vance's meetings with Gromyko, CYRUS first uses context-instantiation strategies to infer that Vance and Gromyko must have talked at diplomatic meetings. Thus, it confines its search to diplomatic meetings. It then searches for diplomatic meetings with Gromyko, and applies strategies to search for larger episodes the meetings could have been part of. See Fig. 21 for the output from CYRUS.

```

Enter next question:
>What have you talked about with Gromyko?

The question is:
((ACTOR HUM1 < = > (*MTRANS*) MOBJECT
  (*CONCEPTS* CONCERNING (*?*)) TO HUM66) TIME G1548)
The question type is "concept completion"

The question concept is:
((ACTOR HUM1 < = > (*MTRANS*) MOBJECT
  (*CONCEPTS* CONCERNING (*?*)) TO HUM66) TIME G1548)
  applying INFER-FROM-PARTICIPANTS to HUM66
    inferring a diplomatic meeting

The inferred question concept is:
((< = > ($MEET ACTOR HUM1 OTHERS HUM66 TOPIC (*?*))
  TIME G1548)
searching memory for question concept
  searching directly for input -- $MEET
  found (GN453)

applying strategies to search memory
  applying FIND-FROM-SIMPLE-MOPS to search for
  episodes $MEET could have occurred in
    searching for sM-SUMMIT-CONFERENCE
    searching for sM-CONFERENCE
    searching for sM-VIPVISIT
    found (GN540 GN481)
    searching sM-VIPVISIT instance for input
    found (GN561 GN564 GN567)
    searching sM-VIPVISIT instance for input
    found (GN485 GN488 GN489)

  applying FIND-FROM-IMOPS to search for episodes
  $MEET could have occurred in
    searching for I-NEGOTIATE
    found (GN391A)
    searching I-NEGOTIATE instance for input
    found (GN542 GN594)

SALT and other arms limitations topics.

```

Figure 21

Similarly, when searching memory to find out whom Vance has discussed SALT with, CYRUS first chooses a context of political meetings, and then applies executive strategies to find them, as shown in Fig. 22.

Enter next question:

>Who have you talked to about SALT?

```

The question is:
((ACTOR HUM1 < = > (*MTRANS*) OBJECT
  (*CONCEPTS* CONCERNING CNTRCT2) TO (*?*)) TIME G1062)
The question type is "concept completion"
The question concept is:
((ACTOR HUM1 < = > (*MTRANS*) OBJECT
  (*CONCEPTS* CONCERNING CNTRCT2) TO (*?*)) TIME G1062)
  applying INFER-FROM-TOPIC to CNTRCT2
    inferring undifferentiated political meeting
The inferred question concept is:
((< = > (sM-MEETING ACTOR HUM1 OTHERS (*?) TOPIC CNTRCT2))
  TIME G1062)
searching memory for question concept
  searching directly for input -- sM-MEETING
    found (GN545 GN525 GN517 GN486 GN453)
  applying strategies to search memory
    checking locative preconditions of input
      could have happened during business trip to USSR
      searching for business trips to POL15
        found (GN481 GN390)
          searching sM-VIPVISIT instance for input
            found (GN485 GN488 GN486 GN489)
          searching sM-VIPVISIT instance for input
        applying FIND-FROM-IMOPS to search for episodes
        sM-MEETING could have occurred in
          searching for I-NEGOTIATE
            found (GN391A GN486)
          searching I-NEGOTIATE for input
            found (GN525 GN545 GN488)
        applying FIND-FROM-SIMPLE-MOPS to search for
        episodes sM-MEETING could have occurred in
          searching for sM-CONFERENCE
            found (GN448)
          searching sM-CONFERENCE instance for input
            found (GN450 GN453 GN456)
        applying FIND-FROM-STANDARDIZATIONS to search for
        standard types of sM-MEETINGS
          searching for $MEET
          searching for $CONSULT
          searching for $PUB-REL-MEET
            found (GN590 GN556 GN561)

```

Carter, Brezhnev, Gromyko, other American and Russian diplomats, and Egyptian vice president Mustafa Khalil.

Figure 22

Because CYRUS's memory organization provides multiple references to the same event, CYRUS often finds the same answer by following multiple paths. In the example described, some of the same meetings were retrieved by searching for trips that were retrieved by searching for negotiations.

5. DISCUSSION

This research is certainly not the first research to claim that human long-term memory is reconstructive. Psychologists as far back as Bartlett (1932) have described memory as reconstructive. More recently, Norman and Bobrow (1975, 1977) have presented a more detailed theory of reconstructive retrieval. A number of recent experimental results illustrate the reconstructive nature of human memory. In asking people to name persons in their high school classes, Williams (1978) found that his subjects recalled features of a person's name such as its first letter, the number of syllables, and what it rhymed with, and used that information to come up with the name. When he asked people to recall whether a karate expert in a particular story had broken a block, Spiro (1979) found that they used their knowledge about karate experts in general to answer the question, rather than retrieving actual story details. In both of these cases only partial information was retrieved from memory, and using generalized knowledge, "actual" items were reconstructed.

In many ways, the work presented here is complementary to the psychological work. While they have proposed general mechanisms for retrieval without worrying about an underlying memory organization or a core retrieval process, this discussion has proposed a memory organization which supports reconstructive processing, and has presented well-specified processes for retrieval which depend on that underlying organization of information and knowledge in memory. The memory organization described is one where items are arranged in memory in conceptual categories, and discriminated within those categories according to their differentiating features. Retrieval is a process of construction and elaboration of contexts for search. This section will discuss the implications of the theory.

5.1 Forgetting

Perhaps the most important observation to make about the reconstructive retrieval process described is that it is fallible. There is always the chance that the right memory will not be remembered. This is true for two reasons. First, the traversal algorithm requires that indices be specified before they are traversed. Thus an item cannot be retrieved unless a suitable discriminating feature can be generated for it. Strategies geared toward specifying plausible features cannot be used to generate deviant features. Thus, items whose features are different, but related to the norms will be easier to find than those whose features deviate excessively. The best example of this comes from examining one of CYRUS's mistakes. During memory update, CYRUS had incorrectly generalized that "diplomatic meetings are about Arab-Israeli peace." That generalization (in the content frame of "diplo-

matic meetings'') had the following effect on retrieval. When asked to enumerate the diplomatic meetings Vance had attended, it used its knowledge about the topics of diplomatic meetings as a constraint on the elaboration it was doing. Thus, it searched for the meetings with features that would have been plausible if the topic were "Arab-Israeli peace," but did not recall other meetings (e.g., meetings about SALT). Its retrieval strategies were working correctly at that point. The knowledge in memory which constrained their application to only relevant features also kept CYRUS from remembering items with deviant features (in this case, an alternate topic). Memory's constraints, then, can cause forgetting.

In addition, the information available in a retrieval specification constrains retrieval. Strategies are dependent on the information or cues available at the time of retrieval. This includes cues available from previous context and those present in a question. In choosing a context for search, for example, the proper context can only be chosen if some aspect of the retrieval specification refers to that context. If a person were answering, "Who is the most famous person you've ever met?", for example, we would not be surprised if he failed to remember meeting a famous person in a museum, since there is no obvious cue to initiate retrieval of museum experiences. On the other hand, if the person were prompted with "How about in a museum?" or if previous conversation had concerned museum experiences, there would be a better chance of his retrieving that experience.

CYRUS also fails in this way. Vance's discussions about the Camp David Accords outside of political meetings, for example, are not recalled by CYRUS in answering "Who have you talked to about the Camp David Accords?" He may have given a speech about the situation or have talked to somebody about the accords at some social political situation (such as a state dinner). Similarly, if Vance's wife had met Mrs. Begin in some situation other than a social activity when the Vance's and Begin's were present, then retrieval would have failed.

The same disadvantage appears in alternate context search. If, during processing, an executive strategy which would yield an answer is neglected, episodes which would have been found by that strategy might not be found. In addition, alternate context search allows retrieval of only standard types of related episodes.

CYRUS is fallible in this regard in trying to recall ceremonies. Because a ceremony commemorating the opening of air traffic routes by flying between the two airports is not a standard type of ceremony, alternate context search is not sufficient for CYRUS to recall that ceremony. If flying had been a standard sort of ceremony, that episode could have been recalled by recalling flying ceremonies. Flying, however, is not standard, so the ceremony could not be recalled in that way.

Failure to apply a relevant strategy can result in failure to retrieve relevant information from memory. Although "Find-from-Standardizations"

is not successful in finding this ceremony, another strategy might work. In some cases, the reason for an event can be useful in recalling it. Since opening air routes is not a standard reason for a diplomatic ceremony, however, that also would not work in this case. Alternatively, recalling ceremonies involved with negotiations might aid retrieval of this particular ceremony.

Another disadvantage of this method of retrieval is that memory is too sensitive to incorrectly specified questions. Partial matching is not allowed. A question may get details of an event wrong. For example, a question could ask about an event that happened in the spring, when in fact the event happened in the summer, or it could ask for a meeting in Jerusalem that actually took place in Tel Aviv. As people, we are often able to recognize and point out another person's mistaken presuppositions and retrieve what was "really" requested. In the scheme described, a mistaken presupposition will keep memory from returning a correct answer.

It should be possible to incorporate a solution to this problem into the retrieval algorithm, although it is not there now. In the retrieval scheme described, memory finds partial matches during its search, but discards those that are inconsistent with the question concept. An evaluation procedure added to the search process might be able to save those matches that are close. In general, however, the problem of partial matching and finding the best partial match has not been solved.

5.2 Directed Retrieval

Why, if there is a possibility of forgetting, is reconstructive retrieval advantageous? The major advantage of reconstructive retrieval is that it is a way of directing search only to relevant places in memory. Component-to-context instantiation strategies, for example, constrain search only to relevant memory categories—in CYRUS's case, relevant E-MOPs. The traversal/elaboration process, along with the indexing scheme, restricts search within memory categories or E-MOPs. If an E-MOP had hundreds or thousands of events, only those directly related to the targeted event would be traversed. Extra effort is put forth in directing search so that the search itself is highly constrained.

5.3 Retrieving Generalized Information

Another feature of this retrieval algorithm and memory organization is that generalized information is retrieved in exactly the same way as individual items. Thus, retrieval for a question such as "When you go to Europe, where do you usually stay?" will proceed exactly as for "Last time you were in Europe, where did you stay?" To answer the second question, the particular trip to Europe would be retrieved by the traversal process (as pre-

viously described), and its context would then be searched for where the actor stayed. In traversing memory to answer the first question, an E-MOP corresponding to "trips to Europe" would be returned by the traversal process in exactly the same way a particular event is retrieved. Generalized information associated with that MOP would then be used to answer the question. A conclusion to be drawn from this observation is that episodic and semantic memory are not separate entities.

5.4 Nonenumeration

There are other conclusions to be drawn from this research, the most important of which is the following:

A few powerful strategies for search, combined with rich conceptual indexing, allow retrieval without category enumeration.

Two types of retrieval strategies have been presented: instantiation strategies, which construct and elaborate contexts for search, and executive strategies, which guide search for and within alternate contexts. In addition, strategy application depends on three types of knowledge being present in memory:

1. context to context relationships
2. component to component relationships
3. relationships between types of components and contexts for search

In order to retrieve events from a category, memory's organization must provide *discriminability* between events. In a memory with that organization, specification of a unique set of features will allow an item to be retrieved. Thus, although retrieval does not involve category enumeration, a series of similar events can be found in a directed way, by reconstructing possible details and searching for events with those details, or by searching for an alternate context that might refer to a target event.

5.5 Retrieval Failures and Ease of Retrieval

The theory presented makes the following psychologically relevant prediction:

The closer a retrieval cue is to the encoding of an item at the time of retrieval, the easier retrieval of that item should be.

An item can be retrieved only by specifying features on which it was discriminated during processing. When an event specification refers to features that have not been discriminated, retrieval can only occur if the specified features can be transformed into features that have been discriminated. Psychologists have called this phenomenon *encoding specificity* (Tulving, 1972). According to this theory, the ease with which an item can be recalled from memory depends on the nearness of the retrieval specification to the description initially encoded in memory for the item.

Memory, however, is constantly reorganizing itself over time. Thus, it is the indices for an event at the time of retrieval and not at time of input which determine ease of retrieval. Any time a retrieval specification refers to features of events which have not been discriminated, it will be necessary to elaborate or reconstruct aspects of the retrieval specification to correspond to features that are discriminated in the conceptual category.

Ease of retrieval depends on ease of finding unique indices: retrieval failure occurs if necessary cues are not available for construction of an appropriate retrieval context. Because reconstructive retrieval processes are knowledge-based, if the knowledge necessary to construct or elaborate a context appropriately is not available at the time of reconstruction, then retrieval will fail.

5.6 Recall and Recognition

Psychologists have often noted that there are many items in memory that can be recognized but cannot be recalled. The CYRUS retrieval processes can be used to explain this. We assume that recognition probes have many more features of an event specified than do recall probes. Because a lot of information is given, there are multiple ways of searching memory without the need for extensive elaboration. Recall probes, on the other hand, provide fewer ways to search memory, and elaboration will be needed. Free recall will require more extensive elaboration than cued recall. Recognition is easier for two reasons: First, it offers more ways to search memory. Second, it allows more directed search of memory without the need to generate features to help direct the search.

5.7 Is CYRUS a Good Model of People?

In previous sections, we have discussed CYRUS's successes as a human model. In evaluating CYRUS's generalizability to human memory, we must also consider features of human memory which CYRUS does not directly address.

First, CYRUS depends on rich indexing of events in memory. Currently, in the implementation each event is indexed by approximately ten

features. Most real-life events that we experience have many hundreds or thousands of features. Thus, we must consider what would happen in a CYRUS-like memory which had to consider that many event features. Would the indexing get out of hand? What is needed to answer that question is the theory of "important feature selection." In part, CYRUS has such a theory, presented in detail in Kolodner (1983, 1984). Because individual events are indexed in many different conceptual structures in memory (e.g., a particular party may be both a political event and a social event), only those features which are domain-related are used as indices in each structure. Second, because indices are only as good as the content-frame of the E-MOP they index, only those features which can make predictions about the probable values of other features make good indices. Thus, we may never have to worry about too many indices. Experimentation based on this model to find out how people index events will shed light on this problem.

Another consideration in deciding if CYRUS is a good model of people is the way it forgets. While CYRUS's errors of omission are obvious, its errors of commission are not. CYRUS never really comes up with a wrong answer. In principle, it should be able to. In implementation, there is no mechanism for it. This is because CYRUS's algorithm combines what is usually thought of as reconstruction (i.e., coming up with an event description) with verification of that description. If it cannot verify a memory, it will not produce it as an answer. CYRUS considers many plausible but wrong memories as it searches memory, but suppresses all of them. It needs a way of judging the veracity of a plausible memory without explicitly "finding" it.

There are a number of features of human memory which we did not consider in developing CYRUS. First, people seem to know what they know and do not know. CYRUS does not. Thus, a person can answer a silly question quickly, make judgments about whether he has remembered all instances of a type of event, make judgments as to whether further memory search will help, etc. CYRUS cannot do any of these things, not out of principle, but because they were not considered in developing the model. Second, CYRUS cannot make recency and frequency judgments. Again, this is a problem we did not consider in developing our framework, but is an important consideration for furthering the theory. Similarly, we did not consider the role of practice in memory. Nor did we consider how incorrect information in a cue affects the retrieval process.

Development of CYRUS was an iterative process of observing people, coming up with a model to explain some behavior, evaluating that model through implementation on the computer, adding features to the model or changing it, then going back and looking at people, etc. Our initial aim was to address the broad problem of reconstructive retrieval from very long term memory and to design a mechanism to explain it. CYRUS is the result. We have been able to provide an algorithm for reconstructive retrieval.

Some aspects of memory which we did not initially consider we got for free from that algorithm—forgetting, generalization, and the combination of episodic and generic memory, to name a few. Other aspects of memory, such as those addressed in the previous paragraph, did not come out of the model for free. Now is the time in the iterative process to consider those problems. We test CYRUS as a model of human memory by running experiments on people to see if its predictions hold and by seeing if additional aspects of human memory can be made to fit *naturally* into the already existing framework.

5.8 CYRUS and EPAM

Some of CYRUS's features are reminiscent of EPAM (Feigenbaum, 1963), an early AI endeavor which modeled human reconstructive memory phenomena. It would be appropriate at this point to compare the two. EPAM's discrimination nets provided an organization, encoding strategy, and retrieval process which modeled the results of human verbal learning experiments. Like CYRUS, EPAM placed new items in its network based on discriminating features, and remembered items by traversing its networks. Like CYRUS, EPAM's processes explained forgetting and false retrieval found in people. There are a number of important differences between EPAM and CYRUS in their organization, encoding, and retrieval processes.

In terms of memory organization and encoding, EPAM's arcs hold negative rather than positive features, providing no place to store generalized information. In addition, only one discrimination is made at each level of the network in EPAM, while in CYRUS, multiple discriminations are made at each level. The ordering of features for discrimination is held constant across items.

There are a number of consequences for retrieval arising from this organization. In discrimination nets, search is guided by the structure of the network rather than the structure of the target item. Questions with no bearing on a target item will be asked of it in the course of memory traversal. An E-MOP's structure, on the other hand, allows retrieval to be directed by the retrieval key or target event itself. Second, EPAM provides only one path to each item. Because multiple discriminations are made at each level in CYRUS, there are multiple paths to each item in memory, allowing it to be retrieved through many different descriptions. CYRUS also provides context-driven strategies for elaborating descriptions and constructing alternate descriptions of items to be remembered.

Finally, it should also be pointed out that the items CYRUS stores (events) are significantly more complex and detailed than those EPAM worked with (nonsense syllables). Thus, CYRUS has a great deal more contextual information available both to use in storing items in memory and in

retrieving them. The use and organization of such knowledge was an important part of deriving CYRUS's structure and processes. While EPAM's organization and processes may be sufficient for explaining people's behavior in verbal learning experiments, they are not sufficient to describe the more complex phenomenon we observe in examining retrieval of episodic information from very long-term memory.

6. SUMMARY

This study has presented a model of reconstructive retrieval, including a memory organization for very long-term event memory, a core retrieval process for searching that memory organization, and a set of retrieval strategies which derive from and correct failures in that core process. This view of remembering accounts for many of the successes and failures of human memory, and also provides a basis for directed retrieval from a large knowledge base. CYRUS implements the retrieval algorithm and memory organization and is able to answer questions by using reconstructive strategies to direct its search. Component-to-component instantiation strategies allow it to narrow its search space by choosing only relevant contexts for search. Component-instantiation strategies allow it to direct search within conceptual categories. Executive search strategies constrain and direct widening of the search space.

Of course, the model is not complete. In addition to the psychological issues mentioned earlier, retrieval strategy application must be controlled, organizations and strategies related to time must be better developed, and reconstructive retrieval of implicit items using retrieval strategies must also be better developed. Work is currently being done on some of these problems (Kolodner, 1981b). So far, we have not had to significantly alter the framework set up here. Rather, we have added to it. In addition, work is being done on testing the psychological validity of CYRUS's strategies, assumptions, and implications (Kolodner & Barsalou, 1982, Reiser, Black, & Abelson, 1982).

REFERENCES

- Anderson, J. R. Retrieval of propositional information from a long term memory. *Cognitive Psychology*, 1974, 5, 451-474.
- Anderson, J. R. *Language, memory, and thought*. Hillsdale, NJ: Erlbaum, 1976.
- Bartlett, R. *Remembering: A study of experimental and social psychology*. London: Cambridge University Press, 1932.
- Bobrow, D. G., & Norman, D. A. Some principles of memory schemata. In D. G. Bobrow & A. Collins, (Eds.), *Representation and understanding*. New York: Academic Press, 1975.

- Cullingford, R. Script application: Computer understanding of newspaper stories. Ph.D. thesis. Research Report #116. Department of Computer Science, Yale University, New Haven, CT, 1978.
- Feigenbaum, E. A. The simulation of verbal learning behavior. In E. A. Feigenbaum & J. Feldman, (Eds.), *Computers and thought*. New York: McGraw-Hill, 1963.
- Kolodner, J. L. Memory organization for natural language database inquiry. Research Report #142. Department of Computer Science. Yale University, New Haven, CT, 1978.
- Kolodner, J. L. Organization and retrieval in a conceptual memory for events. In Proceedings of the International Joint Conference on Artificial Intelligence, Vancouver, BC, Canada, 1981. (a)
- Kolodner, J. L. On beyond CYRUS: Problems associated with long-term memory organization and retrieval. Research Report GIT-ICS-81/17. School of Information and Computer Science, Georgia Institute of Technology, Atlanta, GA, 1981. (b)
- Kolodner, J. L. Maintaining organization in a dynamic long-term memory. *Cognitive Science*, 1983, 7(4). 243-280
- Kolodner, J. L. *Retrieval and organizational strategies in conceptual memory: A computer model*. Hillsdale, NJ: Erlbaum, 1984.
- Kolodner, J. L., & Barsalou, L. W. Psychological issues raised by an AI model of reconstructive memory. Proceedings of the Fourth Annual Conference on Cognitive Science. Ann Arbor, MI, 1982.
- Lebowitz, M. Generalization and memory in an integrated understanding system. Ph.D. Thesis. Research Report #186. Department of Computer Science. Yale University, New Haven, CT, 1980.
- Norman, D. A., & Bobrow, D. G. Descriptions: A basis for memory acquisition and retrieval. Report #7703. Center for Human Information Processing, La Jolla, CA, 1977.
- Reder, L. M., & Anderson, J. R. A partial resolution of the paradox of interference: The role of integrating knowledge. *Cognitive Psychology*, 1980, 12, 447-472.
- Reiser, B. J., Black, J. B., & Abelson, R. P. Retrieving memories of personal experiences. Proceedings of the Fourth Annual Conference of the Cognitive Science Society, Ann Arbor, MI, 1982.
- Schank, R. C. Language and memory. *Cognitive Science*, 1980, 4 (3).
- Schank, R. C., & Abelson, R. P. *Scripts, plans, goals, and understanding*. Hillsdale, NJ: Erlbaum, 1977.
- Schank, R. C., & Kolodner, J. L. Retrieving information from an episodic memory, Research Report #159. Department of Computer Science, Yale University, CT. Short version in Proceedings of the Sixth International Joint Conference on Artificial Intelligence, Tokyo, 1979.
- Smith, E. E., Adams, N., & Schorr, D. Fact retrieval and the paradox of interference. *Cognitive Psychology*, 1978, 10, 438-464.
- Spiro, R. J. Prior knowledge and story processing: Integration, selection, and variation. Technical Report #138. Center for the Study of Reading, IL, 1979.
- Tulving, E. Episodic and semantic memory. In E. Tulving & W. Donaldson (Eds.), *Organization of memory*. New York: Academic Press, 1972.
- Williams, M. D. The process of retrieval from very long-term memory. Center for Human Information Processing, Memo CHIP-75, La Jolla, CA, 1978.
- Williams, M. W., & Hollan, J. D. The Process of retrieval from very long-term memory. *Cognitive Science*, 1981, 5, 87-119.