ORGANIZING MEMORY AND KEEPING IT ORGANIZED

Janet L. Kolodner

Dept. of Computer Science
Yale University, P. O. Box 2158
New Haven, CT 06520

ABSTRACT

Maintaining good memory organization is important in large memory systems. This paper presents a scheme for automatically reorganizing event information in memory. The processes are implemented in a computer program called CYRUS.

INTRODUCTION

People are quite good at retrieving episodes from their long term memories. In fact, they are much better at information retrieval than any current computer system. Psychologists have described human memory as a reconstructive process (e.g., [3]). When people attempt to remember events and episodes from their lives, they often must go through a complicated reasoning and search process ([6] and [5]). These processes are dependent on good memory organization.

In order to keep memory well organized as new data is added, memory organization must support the creation of new memory categories and the building up of generalized knowledge. If a memory held only 10 events that could be described as meetings, a "meetings" category would be useful. But, unless new meeting sub-categories were created as additional meetings were added to the memory, retrieval of meetings would become very inefficient. Thus, a memory system needs the ability to create new categories automatically from old ones.

CYRUS is a computer program which implements a theory of human memory organization and retrieval. The program is designed to store and retrieve episodic information about important people, and is based on a theory of the way people organize and remember information about themselves. Right now, it holds information about Secretaries of State Cyrus Vance and Edmund Muskie. CYRUS answers questions posed to it in English, using search strategies [5] to search its memory and a set of constructive strategies [2] to construct search keys. CYRUS is connected to the FRUMP program [1], which produces summaries of stories off the UPI wire. When FRUMP sends CYRUS new information about Vance or Muskie, CYRUS automatically updates its memory. As it does that updating, it reorganizes its memory and builds up generalized knowledge about the information it stores.

CYRUS' memory is organized around episodes using Memory Organization Packets (MOPs) [4]. Because episodes include references to persons who participated in them, their locations, other episodes they are related to, etc., they are good global organizers of otherwise disjoint information. For example, Camp David and Menachim Begin have in common that Camp David was the location of an important summit conference that Begin attended.

There are a number of problems that must be addressed in maintaining a self-updating memory.

1. what constitutes a good category in memory, i.e., what attributes does a good category have?
2. what kind of knowledge must be stored about each category to enable retrieval and new category creation?
3. how do categories relate to each other?
4. when is it appropriate to reorganize a category into smaller pieces?
5. how can generalized knowledge be added to new categories?

The remainder of this paper will address some of these problems.

RECOGNIZING SIMILARITIES BETWEEN EPISODES

People notice common aspects of episodes and make generalizations in the normal course of understanding. Reorganization of memory requires noticing similarities and making generalizations based on those similarities. Generalized knowledge is needed to predict future occurrences, to elaborate on a context being understood, to help direct memory update, and as an aid in directing search during retrieval. Like people, CYRUS notices similarities between episodes and makes generalizations from them.

Similar episodes in CYRUS are stored in the same MOP, along with the generalized knowledge built up from them. MOPs act as event categories in memory holding episodes and knowledge about those episodes. The generalized information a MOP

---

*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.

holds resides in its "content frame" [4] and includes such things as typical preconditions and enablement conditions for its episodes, their typical sequence of events, larger episodes they are usually part of, their usual results, typical location, duration, participants, etc.

The structure of individual episodes provides a framework for deciding whether two episodes are similar to each other. If, on two different trips, Vance is welcomed at the airport and then has a meeting with the president of the country he is in, then the episodic structure of the two trips will look alike, and we can say that the two episodes are similar. While on the second trip, he might be reminded [4] of the first one because of their similarities. In the same way, a person hearing about the two trips might be reminded of the first when hearing about the second. If the result of the first trip had been a signed accord, then he may predict, or at least hope, that an accord would be signed at the end of this trip also. If an accord is signed, he will generalize that when the first event of a diplomatic trip is a meeting with the head of state, then an accord will be reached. Later, he will be able to use that knowledge in understanding and retrieval.

REORGANIZING EVENT CATEGORIES

In order for such reminding and subsequent generalization to occur in CYRUS, its MOPs must be highly structured internally. Episodes are indexed in MOPs according to their important aspects. Indexing in a MOP is by content-frame components and includes sequence of events, location, participants, duration, etc. When an event is indexed similarly to an event already in memory, reminding occurs and generalizations are made based on their similarities. As a result, a sub-MOP of the larger MOP is formed to hold those episodes and their generalizations. With the addition of more episodes, new sub-MOPs are further divided. In creating new MOPs and building up generalized information, new knowledge, which can be used for later understanding and retrieval, is added to the data base.

The actual processing when an event is added to a MOP depends on its relationship to events already in the MOP. One of the following four things is true about each component of an event description:

1. It is unique to this event in the MOP
2. It is semi-unique to this event in the MOP (it has happened once or a small number of times before)
3. It is often true of events in the MOP
4. It is typical of events in the MOP

In case 1, when the descriptive property is unique, the event is indexed under that aspect in the MOP. For instance, one of the discriminations CYRUS makes on meetings is the topic of the contract being discussed. A meeting about the Camp David Accords is indexed under "contract topic = peace", and a meeting about military aid to Pakistan is indexed under "contract topic = military aid". The first time CYRUS hears about a meeting in which Vance discusses military aid, it will index that meeting uniquely in the "diplomatic meetings" MOP under the property "contract topic = military aid". If it were also the first meeting he had with a defense minister, then it would also be indexed uniquely under the occupation of its participants (because meetings are occupational).

If a property being indexed has occured once before in a MOP (case 2), then reminding occurs, the two events are compared to see which other common aspects they have, and generalizations are made. When Vance meets again about military aid, CYRUS is reminded of the prior meeting because both have the same topic. It checks the descriptions of both to see what other similarities they have. If both, for example, are with defense ministers, it will conclude that meetings about military aid are usually with defense ministers. It also begins indexing within the new MOP:

```
*************
Adding
   $MEET actor (Vance)
         others (defense minister of Israel)
         topic (Military aid to Israel)
         place (Jerusalem)
to memory ...

Reminded of
   $MEET actor (Vance)
         others (defense minister of Pakistan)
         topic (Military aid to Pakistan)
         place (Washington)

because both are "diplomatic meetings"
        both have contract topic "military aid"
    creating new MOP: meetings about military aid
  generalizing that when
    Vance meets about military aid,
    often he meets with a defense minister
*************
```

Later, if CYRUS hears about a third meeting whose topic is military aid, it will assume that the meeting is with the defense minister of the country requesting aid (unless it is given contrary information). If asked for the participants of that event, it will be able to answer "probably the defense minister". If, on the other hand, a number of meetings about military aid with participants something other than defense ministers are added to memory, CYRUS will remove that generalization and attempt a better one instead.

On entering the next meetings about military aid to memory, CYRUS will index them along with other events already indexed there. A new meeting about military aid will be entered into the "meetings about military aid" sub-MOP of "diplomatic meetings" and will be indexed within that MOP (case 3). In this way, reminding, generalization, and new MOP creation will occur within newly created MOPs. If a new meeting about military aid to Pakistan is added to memory, CYRUS will be reminded of the first because both will be indexed under "contract sides = Pakistan" in the "meetings about military aid" MOP.

No discrimination is done on properties that are typical (case 4) of events in a MOP (i.e., almost all events in the MOP fit that description). In that way, generalization can control the expansion of MOPs in memory. If memory has generalized that meetings are called to discuss contracts, then the fact that the topic of a later meeting is a contract will never be indexed. Appropriate aspects of the contract, however, will be indexed.

Thus, if a new event with no unique aspects is added to memory, reminding of individual events does not occur, but generalizations already made are confirmed or disconfirmed. Disconfirmed generalizations are removed. When CYRUS hears about yet another meeting in the Mid East about the Camp David Accords, it will not be reminded of any specific episodes, but will put the new meeting into the MOPs it fits into.

```
**********
Adding
   $MEET actor (Vance)
         others (Begin)
         topic (Camp David Accords)
         place (Jerusalem)
to memory ...

Putting it into MOP:
       meetings with Begin in Israel
    confirming generalizations
Putting it into MOP:
       meetings about the Camp David Accords
                with Israeli participants
    confirming generalizations
Putting it into MOP:
       meetings in Israel
    confirming generalizations
         ...
**********
```

IMPLICATIONS IN RETRIEVAL

What are the implications of this indexing scheme in retrieval? The retrievability of an event depends on how distinct its description is, or how many of its features turn out to be significant. As events with similar properties are added to memory, their common aspects lose significance as good retrieval cues and category specifiers (case 4). An event with no unique or semi-unique descriptors will become lost in memory or "forgotten". Since events are indexed by their differences, they can be retrieved whenever an appropriate set of those differences is specified, but specification of only common aspects of events will not allow distinct episodes to be retrieved.

Generalized knowledge can be used during retrieval in a number of ways. One important use is in guiding search strategy application (see [5]). Generalized knowledge can also be used for search key elaboration. There is not always enough information given in a question to direct search to a relevant MOP or to a unique episode within a MOP. Generalizations and a MOP's indexing scheme can be used to direct the filling in of missing details. Only those aspects of a MOP that are indexed need be elaborated.

Generalized information can be used to answer questions when more specific information can't be found. If CYRUS has made a generalization that Gromyko is usually a participant in SALT negotiations, it will be able to give the answer "probably Gromyko" to "Last time Vance negotiated SALT, who did he meet with?", even if it could not retrieve the last negotiating episode. In the same way, generalizations can be used for making predictions during understanding.

CONCLUSIONS

Good memory organization is crucial in a large information system. Some important processes memory organization must support include dynamic reorganization of memory, creation of new memory categories, and generalization. In this paper, we've tried to show how a long-term memory for episodic information can keep itself usefully organized. This requires a good initial organization plus rules for reorganization. Reorganization of memory happens through reminding, or noticing similarities between episodes, and generalization. The generalizations produced are useful both for controlling later memory reorganization and for retrieval. Some related problems which have not been addressed in this paper, but which are important, are updating generalized knowledge (particularly recovery from bad generalizations), judging the usefulness of new categories, and guiding indexing properly so that new MOPs and generalizations are useful and relevant. Those topics are currently being addressed.

REFERENCES

[1] DeJong, G. F. (1979). Skimming stories in real time: An experiment in integrated understanding. Research Report #158. Department of Computer Science, Yale University.

[2] Kolodner, J. L. (1980). Organization and Retrieval from Long Term Episodic Memory. Ph.D. Thesis (forthcoming). Department of Computer Science, Yale University.

[3] Norman, D. A. and Bobrow, D. G. (1977). Descriptions: a basis for memory acquisition and retrieval. Report #7703. Center for Human Information Processing, La Jolla, California.

[4] Schank, R. C. (1979). Reminding and memory organization: An introduction to MOPs. Research Report #170. Department of Computer Science, Yale University.

[5] Schank, R. C. and Kolodner, J. L. (1979). Retrieving Information from an Episodic Memory Research Report #159. Department of Computer Science, Yale University. also in IJCAI-6.

[6] Williams, M. D. (1978). The process of retrieval from very long term memory. Center for Human Information Processing Memo CHIP-75, La Jolla, California.