

RETRIEVING INFORMATION FROM AN EPISODIC MEMORY
OR
WHY COMPUTERS' MEMORIES SHOULD BE MORE LIKE PEOPLE'S *

Roger C. Schank and Janet L. Kolodner
Computer Science Department
Yale University
Box 2158 Yale Station
New Haven, CT 06520

A computer memory organization modeled after human memory is proposed, along with strategies for accessing that memory. CYRUS, a program which implements the theory, is described.

Do A.I. programs understand what they read? Much debate has taken place on this issue (e.g., [6]) with little agreement. One thing seems clear, however. No computer program can seriously be said to have understood anything unless it can pass a simple test of remembering what it has supposedly understood. This problem seems obvious and easily solved. We can point to programs that paraphrase or summarize. Clearly, they remember. Or do they?

Memory, after all, also involves the integration of new knowledge with old knowledge and the ability to use newly obtained information for future and as yet unspecified tasks. In our own laboratory, we have developed a number of story understanding programs that summarize or answer questions about the stories they have understood (e.g. SAM, PAM, FRUMP). But none of those programs can apply information gained in reading one story to understand a later story. None have a long-term memory.

What would a long-term memory look like? Psychologists have proposed a number of theories, none of which are completely adequate for describing the structure of memory (e.g. [3], [7], [1]). One way to discover what human memory looks like, and what computer memory should look like, is to look at the memory tasks humans perform that we would want a computer to do.

Some of the memory tasks that people do in the course of normal conversation and understanding are listed below (see [5] for more detail):

*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.

1. Reminding
2. Relating events by time.
3. Recognizing when statements or questions are non-sensical.
4. Recognizing when something has not happened.
5. Making inferences.
6. Using time and place specifications.
7. Using maps and other aids.
8. Extensive memory search.

If we want to build an understanding machine that understands in the same way people do, then we have to take into account the vast number of memory tasks that people do. Reminding, for example, is a prerequisite to the process of creativity. Machines that cannot make new, undirected connections between items in memory will never be able to dream or imagine. Machines that cannot be reminded will not be able to think in any serious sense.

To get a handle on the way humans do these tasks, we examined people's responses to a request to tell us about all the museums they had been to ([5]). The results are summarized below:

1. Find museums by searching for instances of the museum script.
2. Find museum experiences (instances of the museum script) by searching for them as events within trip instances.
3. Find museums by searching memory for cities visited.
4. Find museum experiences by searching for experiences in places of residence.
5. Find trips by searching for trips to foreign countries or within country of residence, or by searching for episodes associated with those places.
6. Find places by using a map of the general area.

Rules such as these can be stored in the

computer as features of the scripts and other stereotypes that they refer to. A memory set up this way would have its scripts and other stereotypes pointing to information about where in memory their instances could be found. In this way, each stereotype would specify a set of strategies for searching memory to find its instances. Each script and stereotype would include information about what other scripts or macro-scripts it could be a part of (e.g., the museum script can be an event in a trip macro-script); which of its parts are important cues for finding its instances (e.g., look for places different than home to find trips); and what other memory structures can be used to find its instances (e.g., go through a mental map of an area to find places). Using this organization, the following general rules for searching memory can be used:

7. To find instances of a script, search for other scripts and macro-scripts it can be part of.
8. To find objects, search for stereotypical events (such as scripts) associated with that object.
9. To find an instance of a stereotype in memory, find instances of the parts of it that are specified as important.
10. To find places, use a map.

In order to apply these strategies to a memory search, memory must be well-organized. It should be able to answer "Where did John go to law school?" without having to go through John's experiences in elementary school. It seems that experiences that happen to a person are grouped in terms of the larger contexts in which they were embedded. Knowing when an event might have occurred, or what events might have occurred before or after it, can help us know where in memory to look for it.

Thus, one good way of organizing experiences is temporally, but with the time line broken into workable parts, each with some unique characteristic. These pieces are called eras ([2]). The three years of going to law school, for example, make up a professional school era. Some standard ways of breaking up a personal time line are: places lived, jobs held, schools attended, and family situation. An event in a person's life would be organized on the event list of its corresponding era or eras.

Strategies for searching memory can take advantage of this memory organization. When searching for any particular event, only eras relevant to it must be searched. To find an

event having to do with someone's job, only relevant occupational eras will have to be searched. To implement this on a computer, each script and other stereotype needs to have information about which times in someone's life he might be involved in that type of activity, and which types of eras and era sequences it can be found in.

The following additional rules can be used for searching a memory organized as described above:

11. To find an instance of a stereotype in memory, search the types of eras and era sequences it could be found in.
12. To find an instance of a stereotype in memory, search the appropriate eras going on at the time when this event is most likely to have taken place.
13. To find an era, search in the appropriate era sequence at the time specified relative to other eras in that sequence and in other sequences.

The CYRUS system ([2]) is a memory model that organizes biographical information about people and uses knowledge about its organization for retrieval and automatic updating. An important aspect of the CYRUS system is that the organization of its memory represents an attempt to model memory in people, and its retrieval and updating procedures mirror the way we believe people access their memories. CYRUS is an implementation of the ideas about memory organization and retrieval described above. A sample dialogue with CYRUS, showing input and output, follows:

- Q1: Who is Cyrus Vance?
A1: United States Secretary of State.
Q2: How did he become Secretary of State?
A2: He was appointed by President Carter.
Q3: Where is he today?
A3: Egypt.
Q4: Why did he go there?
A4: He went to negotiate a compromise agreement to Egypt.
Q5: Has he met with the Israelis recently?
A5: Yes, in Washington right before he left for Egypt.
Q6: Who was at the meeting?
A6: Moshe Dayan and Israeli legal experts.
Q7: Will he meet with them in Israel?
A7: Yes, to discuss the compromise.

Each of the memory structures and strategies mentioned above is implemented in CYRUS. The CYRUS episodic memory contains events in the lives of the people represented. CYRUS' stereotype memory holds all the stereotypes

that can be used for representation along with information relevant to each of them.

Eras in CYRUS are characterized by the major role themes a person is involved in at each point in his life ([4]). In CYRUS, each person represented has an occupational sequence of eras (including school), a family sequence of eras, a social sequence of eras, and a residential sequence of eras (where each era corresponds to a place the person has lives). Eras contain all the events in a person's life occurring during that time span and pertaining to that role theme or place of residence. Thus, Cyrus Vance's current occupational era is characterized by his being Secretary of State; his current family era is characterized by his being a husband and father.

CYRUS makes use of information attached to its scripts and other stereotypes to apply the strategy rules described above (rules 7-13). In answering the question, "When was the last time Vance met with Begin?", CYRUS uses information from the attend-meeting script to guide its memory search. CYRUS searches for the meeting by looking in eras and macro-scripts specified by the meeting script as appropriate to meetings. CYRUS' general knowledge tells it that the place for a meeting is usually the country of residence of one of the participants; and that a person must take a trip to get to a place other than his residence. A trip's type and location in memory are determined by the events which happen during the trip, so CYRUS searches for the trip by searching in an era appropriate to the original event being searched for. Input and output from CYRUS follow:

```
>When was Vance's last meeting with Sadat?
  Searching Vance's occupational eras for
  meetings with Sadat
  Searching Vance's occupational eras for
  trips to Egypt
  Found CON562
  Searching CON562 for meetings with Sadat
  Found CON547
  Searching CON547 for meetings with Sadat
  Found one
  December 15, in Egypt.
```

CYRUS uses the strategy of searching for meetings to answer this question. Alternative strategies could have also been used to answer this question. Experiences associated with Sadat could have been searched, and a recent meeting with Sadat might have been found that way. Alternatively, Vance's experiences associated with the Arab-Israeli conflict or the Egyptian-Israeli peace talks could have

been searched (after inferring that a meeting between Vance and Sadat probably had to do with one of those issues).

Eras and era sequences in Cyrus can also be used to relate events through time. If asked "When did Vance get married?", CYRUS would answer "Soon after he began to work as a lawyer." If asked, "When was he last in Paris?", it would answer "During the Vietnam peace talks, when he was an advisor to President Johnson."

CYRUS infers things not explicitly in memory by making use of expectations derived from the scripts, macro-scripts, and role themes it knows about. It can use its knowledge about enablement conditions for role themes to answer questions such as "How did Vance become Secretary of State?" It uses its knowledge about trips to answer questions such as "How did Vance get to Russia?" by inferring that he took an airplane, even if that information is not explicitly in memory.

References

- [1] Bobrow, D. G. and Norman, D. A. (1975). Some principles of memory schemata. In Bobrow, D. G. and Collins, A., ed. Representation and Understanding. Academic Press, Inc., New York.
- [2] Kolodner, J. L. (1978). Memory organization for natural language database inquiry. Research Report #142. Department of Computer Science. Yale University, New Haven, CT.
- [3] Quillian, M. R. (1968). Semantic Memory. In M. Minsky (ed.), Semantic Information Processing. MIT Press, Cambridge, MA.
- [4] Schank, R. C. and Abelson, R. P. (1977). Scripts, Plans, Goals, and Understanding. Lawrence Erlbaum Press, Hillsdale, N.J.
- [5] Schank, R. C. and Kolodner, J. L. (1979). Retrieving information from an episodic memory (long version). Research Report #157. Department of Computer Science. Yale University, New Haven, CT.
- [6] Weizenbaum, J. (1977). Computer Power and Human Reason: From Judgement to Calculation. Freeman, San Francisco.
- [7] Williams, M. D. (1978). The process of retrieval from very long term memory. Center for Human Information Processing Memo CHIP-75. University of California at San Diego.