10

The structure of memory

ANSWERING QUESTIONS

When to retrieve information
Retrieving an image
Retrieval as problem solving
Retrieval and comprehension

THE IMPORTANCE OF UNDERSTANDING

A MODEL OF MEMORY

Remembering concepts
Primary and secondary concepts
Remembering events

SUGGESTED READINGS

Lindsay, Peter H. and Norman, Donald H. **Human Information Processing: An Introduction to Psychology**. New York: Academic Press, 1972. 773 p. — Ch. 10 The structure of memory, p. 374-401; Ch. 11 Memory processes, p. 402-434

The purpose of this reading is not to learn about human memory, but rather to learn about search and information structures that support search, in this case semantic networks and frames. In fact, psychologists today have other models of how memory works. (Any one memory model, including semantic networks, may provide a different angle on the working of memory, but again this is not the point of this reading from a 670 perspective. There is an interchange between artificial intelligence and psychology: Understanding how memory and reasoning work in people provides ideas for intelligent computer systems, and the design of intelligent computer systems provides ideas for models of how the brain works.)

The beginning of Chapter 10 can be read as good advice to reference librarians on how to approach a new search. Semantic networks and frames are used in intelligent computer information retrieval and question-answering systems.

The sections marked ⌈ ⌋ are not essential to the argument and can be omitted.

In the house you lived in two houses ago, as you entered the front door, was the doorknob on the left or right?

Good general introduction to any kind of retrieval

Here is a question that requires memory and the retrieval of information. But if you try to answer it, you will find that the task seems more like that of solving a problem than of retrieving something from memory.

In fact, studies of remembrance, problem solving, thinking, and mental operations have much in common, since there is little to distinguish among them. A person recalling material seems to be solving a problem. First he analyzes the question to decide whether it is legitimate, whether he is likely to have the information, and if so, how difficult it will be to find. If he decides to attempt a recall, he sets up a retrieval strategy. As he proceeds, he combines the information in the request with partial solutions to form new questions and continue his search. His retrieval path seems to be organized around prominent events, landmarks in his memory that stand out above the myriad of stored details. Even if he recovers the information requested, much of his recollection appears to follow from logic and a reconstruction of what must have been.

The study of long-term memory is the study both of this problem-solving process and of the structure of the memory on which it operates. The things we remember are organized into a complex structure that interconnects the events and concepts built up by past experience. The act of remembering is the systematic application of rules to analyze this stored information.

ANSWERING QUESTIONS

Perhaps nowhere else is the power of human memory so clearly exhibited as when people answer questions about what they know. Consider what is required to answer a question (hereafter called a *query*). First, it is not sufficient just to have the pertinent information stored in memory. It is necessary to search out and find all the stored information that is relevant to the query, to evaluate any contradictory data, and finally to put it all together to form the best answer, given the information retrieved.

The human brain is not the only system faced with the problem of answering questions based on large amounts of information. There are numerous examples of systems that are capable of holding vast amounts of data: They range from such traditional devices as libraries to modern computer-based systems. When working with such memories, the first thing that is discovered is that getting information into the system is usually not a basic problem. The difficulties arise in trying to get it out.

Regardless of the memory system we contemplate—be it the human brain, a library card-catalog, a large collection of file folders, or a computer mass-storage device—there are types of queries for which the organizational structures of the memory (including its indices and abstracts) are inappropriate. Yet an omniscient outside observer might declare that the information required to answer the query does, in fact, exist in the system, if only the user would ask the correct questions and then put together the results sensibly. How can the system be designed so that, after we have gone to all the trouble of collecting information, we can find the information we want? What kinds of retrieval strategies are required? If the questions to be asked are known beforehand, the problem may not be too difficult. It is relatively easy, for example, to design the census system so that it can quickly find out how many people are under the age of 30, if you know you want that information before you store the data. But what about questions that were not anticipated? Is it possible to build a data-processing system that, like the human memory, can answer almost anything that comes to mind?

The key to any large-scale memory system, then, is not its physical capacity for storing huge amounts of information. Rather, it lies in its ability to retrieve selected pieces of data on request, its ability to answer questions based on the information stored. We can learn a great deal about the nature of the data-processing operations involved in human memory simply by sitting back and thinking about the kinds of questions people can answer and of the mechanisms and procedures necessary to answer them.

When to retrieve information

Query: **What was Beethoven's telephone number?**

What is your answer to this query? Nonsense, you say. Beethoven died before telephones were invented. But suppose we ask about someone who had a telephone?

Query: **What was Hemingway's telephone number?**

You still refuse to try to retrieve the number. You don't know. How do you know you don't know? What about:

Query: **What is the telephone number of the White House?**

What is the telephone number of your best friend?

What is your telephone number?

The principle being illustrated is that, when asked to recall something, you do not start off blindly on a search. First you seem to analyze the question to see whether you are likely to find anything. On the basis of this preliminary analysis, you may conclude that there is no sense in even attempting to recall the data. Maybe the information does not exist. Maybe the information exists, but you know it is not in your memory. But what information do you use to decide that you do not know Hemingway's phone number, even if he had a phone? Maybe you think you might be able to retrieve the information if you tried, but it would require too much effort to be worth the bother. Are you really sure you could not produce the White House number if you worked on the problem for awhile?[1]

When we ask questions of human memory we discover that there are procedures that analyze the message to determine if the relevant information exists, whether it is likely to have been stored, and the effort required and probable success of an attempt at retrieval. This whole sequence of operations seems to be carried out rapidly and unconsciously. We are only vaguely aware of the complexity of the rules involved.

Clearly, such a system is a great advantage for a large-scale memory. It does not waste time looking for things it does not know. It can judge the cost of retrieving information that is difficult to find. We will see later that when faced with a continuous bombardment of sensory information, it is very important to know what is not known, since it lets us concentrate on the novel, unique, important aspects of events in the environment.

We cannot yet define these preprocessing procedures well enough to take advantage of them in designing a retrieval system. We know such mechanisms operate in human memory. We can identify some of the basic processes and outline their general properties. But we cannot yet describe the details of the machinery involved.

Retrieving an image

Query: **In the rooms you live in, how many windows are there?**

This time, retrieval should proceed smoothly. First you conjure up an image of each room, then examine it, piece by piece, counting the windows. You then move to the next room and continue the process until you finish. The task seems easy. Yet, apart from the fact that people can have and use images, very little is known about the nature of internal images, how they are stored or how they are retrieved.

Whether present theories can handle the problem or not, it is clear that our memories do contain a large number of images of our past

[1] The White House telephone number is (202) 456-1414.

experiences. An image can be retrieved and examined at will: the face of a friend, a scene from our last trip, the experience of riding a bicycle. This record of visual experiences suggests some important principles for the analysis of retrieval strategies. Saving some form of a replica of the original information provides a great deal of flexibility in being able to deal subsequently with questions about experiences. It is unlikely that you thought about the possibility of someday being asked for the number of windows. There is no need for you to take note of this fact whenever you are in your room. As long as you save an image of the rooms, you can worry later about retrieving particular pieces of information when they may be required.

We do not always deal with visual information by storing it all away. Often we analyze and condense incoming information, throwing away irrelevant details and remembering only what seems important. Try to recall what we have said so far in this chapter. You do not conjure up an image of the pages and read off the words. You recall a highly abstract version of your visual experience, reorganized and restated in your own terms.

An adequate model of human memory, then, will have to describe when incoming events are saved in their entirety and when only the critical features are extracted and stored. Recording a replica of the information uses up considerable memory space, makes subsequent retrieval more complicated and time consuming, and tends to clutter up the memory with irrelevant details. Reorganizing and condensing the information to save only the central features runs the risk of failing to record information that might subsequently be important. It limits the range and variety of ways in which past experiences can be used and the types of questions we can answer. Maybe it would be optimal to save both a complete record and a reorganized, condensed version, or maybe there are more sophisticated ways of dealing with rote records. Are there general rules for recording and reconstructing images that simplify the storage problem without sacrificing details? After all, houses have lots of things in common, such as roofs and walls. Perhaps the human memory system capitalizes on these similarities.

Regardless of how the information is actually stored, it is important to have both some form of an image of the rooms and a procedure for counting the windows. During retrieval, these two processes interact: One retrieves and constructs the image; the other analyzes and manipulates the retrieved information. Just as for problem solving, retrieval requires the active construction and analysis of information through the application of rules or procedures. This constructive aspect of human memory comes out more clearly when the system is presented with yet another kind of question.

continue on p. 385

Retrieval as problem solving

Query: **What were you doing on Monday afternoon in the third week of September two years ago?**

Don't give up right away. Take some time to think about it and see if you can come up with the answer. Try writing down your thoughts as you attempt to recover this information. Better still, ask a friend to think out loud as he tries to answer the query.

The type of responses people typically produce when asked this kind of question goes something like this:

1. *Come on. How should I know?* (Experimenter: Just try it, anyhow.)
2. *OK. Let's see: Two years ago. . . .*
3. *I would be in high school in Pittsburgh . . . .*
4. *That would be my senior year.*
5. *Third week in September—that's just after summer—that would be the fall term. . . .*
6. *Let me see. I think I had chemistry lab on Mondays.*
7. *I don't know. I was probably in the chemistry lab. . . .*
8. *Wait a minute—that would be the second week of school. I remember he started off with the atomic table—a big, fancy chart. I thought he was crazy, trying to make us memorize that thing.*
9. *You know, I think I can remember sitting. . . .*

Although this particular protocol is fabricated, it does catch the flavor of how the memory system works on this kind of retrieval problem. First, the question of whether or not to attempt the retrieval: The preliminary analysis suggests it is going to be difficult, if not impossible, to recover the requested information and the subject balks at starting at all (line 1). When he does begin the search, he does not attempt to recall the information directly. He breaks the overall question down into subquestions. He decides first to establish what he was doing two years ago (line 2). Once he has succeeded in answering this question (line 3), he uses the retrieved information to construct and answer a more specific question (line 4). After going as far as he can with the first clue, he returns to picking up more information in the initial query, "September, third week." He then continues with still more specific memories (lines 5 and 6). Most of what happened between lines 7 and 8 is missing from the protocol. He seems to have come to a dead end at line 7, but must have continued to search around for other retrieval strategies. Learning the periodic table seems to have been an important event in his life. The retrieval of this information seems to open up new access routes. By line 8, he once again appears to be on his way to piecing together a picture of what he was doing on a Monday afternoon two years ago.

Here memory appears as a looping, questioning activity. The search is active, constructive. When it cannot go directly from one point to another, the problem is broken up into a series of subproblems or subgoals. For each subproblem, the questions are: Can it be solved; will the solution move me closer to the main goal? When one subproblem gets solved, new ones are defined and the search continues. If successful, the system eventually produces a response, but the response is hardly a simple recall. It is a mixture of logical reconstruction of what must have been experienced with fragmentary recollections of what was in fact experienced.

This idea of memory as a problem-solving process is not a new notion. Similar ideas have been suggested by poets and philosophers for thousands of years. For psychologists, early and persuasive proponents were William James (1890) and Sir Frederick Bartlett (1932). What is new is that finally there are some analytic tools to deal with such processes in detail. The machinery is available to build and test models of memory that solve problems by breaking up questions into subgoals, that try to converge on solutions through continued reformulation and analysis of promising subquestions.

Retrieval and comprehension

Query: **Can pigeons fly airplanes?**

You should be fairly fast on this one. Assuming you responded negatively, then obviously the next question is, why not? Could a pigeon, in principle, fly an airplane?

This time, the problem is not one of preprocessing the message or of analyzing visual images, or of setting up a search strategy to recover specific information. Moreover, the question does not seem to be answered in terms of a simple recall. It is unlikely that you have given much thought to the specific possibility that pigeons might be pilots, or have stored directly the fact that they do not fly planes as part of the information associated with pigeons. Somehow, you arrive at your answer by a logical analysis of the information associated with the two concepts. But you are only vaguely aware of your analysis when you make your first impulsive response. It takes quite a bit of subsequent work to discover exactly why it is unlikely that pigeons fly airplanes. On further thought, you may even reverse your original verdict: Perhaps a pigeon (with a few modifications either to the pigeon or the airplane) could fly an airplane, after all.[2]

[2] When this question was posed to one of the authors' daughters (8 years old), the first response was laughter at the apparently preposterous idea. The second response was that maybe it was not so silly, since both pigeons and pilots do fly.

No existing data-storage system, except the human brain, has the capability of answering this type of question. By contrast, humans routinely analyze the information they receive to test its consistency with their past experience to decide whether what is being said is plausible in terms of what is already known.

This behavior indicates that human memories cannot be a random collection of facts. On the contrary, the information in memory must be highly interrelated and structured. Concepts can easily be compared to deduce the similarities and differences. Moreover, the comparison seems to go on at a number of levels. During the first stage of analysis, something must have suggested immediately that the concepts of pigeons and flying airplanes are logically inconsistent. The second stage requires a more extensive analysis which seems to be similar to the problem-solving procedures just discussed. The subquestion is asked as to what is required to convert a pigeon into a pilot and the stored information is reevaluated on the basis of this reformulation of the question.

This last question brings us to a key issue in the study of memory: the ability to understand. Somehow humans comprehend the meaning of an input and use this comprehension to evaluate present experiences in terms of what is already known. Any model of memory must come to grips with this basic aspect of long-term memory in humans.

THE IMPORTANCE OF UNDERSTANDING

Comprehension is an illusive process. How does one define exactly what it means to understand a word? In fact, how can one even decide from external appearances whether a machine or a person comprehends what is being said? If you think you can, then consider these two dialogs: They have been produced by two different information-processing models of memory which, in these instances, have been translated into computer programs.

The first is a dialog between a person and a program called DOCTOR, a computer program designed to simulate a psychiatrist. The person using the program is allowed to type anything he wishes onto the typewriter keyboard. The second is a program that moves blocks around on a table. The statements of the human user appear in `this typeface and are always preceded by this symbol` ●; the responses of the computer to these statements are printed in this print.

Figures 10-1 & 10-2

Examine the dialogs. Do you think the computers understand what the people are saying? The program conducting the psychiatric interview is a fraud. It does not have any idea at all what **boyfriends** are, or what it means to be **depressed** or **aggressive**. Nevertheless, to the casual observer, it certainly puts up a convincing front. The second program knows quite a bit about blocks and language. But does it really have

FIGURE 10-1

- `Men are all alike.`
  In what way
- `They're always bugging us about something or other.`
  Can you think of a specific example
- `Well, my boyfriend made me come here.`
  Your boyfriend made you come here
- `He says I'm depressed much of the time.`
  I am sorry to hear you are depressed
- `It's true. I am unhappy.`
  Do you think coming here will help you not be unhappy
- `I need some help, that much seems certain.`
  What would it mean to you if you got some help
- `Perhaps I could learn to get along with my mother.`
  Tell me more about your family
- `My mother takes care of me.`
  Who else in your family takes care of you
- `My father.`
  Your father
- `You are like my father in some ways.`
  What resemblance do you see
- `You are not very aggressive but I think you don't want me to notice that.`
  What make you think I am not very aggressive
- `You don't argue with me.`
  Why do you think I don't argue with you
- `You are afraid of me.`
  Does it please you to believe I am afraid of you
- `My father is afraid of everybody.`
  What else comes to mind when you think of your father
- `Bullies.`
  Does that have anything to do with the fact that your boyfriend made you come here. (From Weizenbaum, 1968, pp. 172–173)

FIGURE 10-2

- `Will you please stack up both of the red blocks and either a green cube or a pyramid?`
  Ok
- `Which cube is sitting on the table?`
  The large green one which supports the red pyramid
- `Is there a large block behind a pyramid?`
  Yes, three of them: A large red one, a large green cube and the blue one
- `Put a small one onto the green cube which supports a pyramid.`
  Ok
- `Put the littlest pyramid on top of it.`
  Ok
- `Does the shortest thing the tallest pyramid's support supports support anything green?`
  Yes, the green pyramid
- `What color is it?`
  By "it", I assume you mean the shortest thing the tallest pyramid's support supports
  Red
- `How many things are on top of green cubes?`
  I'm not sure what you mean by "on top of" in the phrase "on top of green cubes"
  Do you mean:
  1—Directly on the surface
  2—Anywhere on top of?
- `Anywhere on top of`
  Three of them
- `Had you touched any pyramid before you put the green one on the little cube?`
  Yes, the green one
- `When did you pick it up?`
  While I was stacking up the red cube, a large red block and a large green cube
- `Why?`
  To get rid of it
- `Why did you do that?`
  To clean off the red cube
- `Why did you clear off that cube?`
  To put it on a large green cube
- `Why did you do that?`
  To stack up the red cube, a large red block and a large green cube
- `Why did you do that?`
  Because you asked me to
- `How did you do it?`
  By putting a large red block on the table; then letting go of it; then putting a large green cube on it; then letting go of that cube; then putting the red cube on that cube; then letting go of that cube
- `How many objects did you touch while you were doing it?`
  Four of them
- `What did the red cube support before you started to clean it off?`
  The green pyramid
- `There were five blocks to the left of the box then.`
  No, only four of them: the red cube, two large green cubes and a large red block
- • •
- `Thank you.`

You're welcome! (From Winograd, 1970, pp. 35–60) the same understanding of the terms as the person with whom it is conversing? How would you test it to find out?

Like the processes of perception, the processes involved in understanding a message go on so quickly and automatically that normally we only become aware of them when they fail. Also like perception, we learn about the memory system by considering the kinds of inputs which are not dealt with smoothly and easily. For example, consider this statement:

Oranz minsocs tankin rakans.

You can read every word. Yet you immediately dismiss the statement as meaningless, and you are right: It is meaningless. But how did you know? Did you really check all the words you have encountered during the many years of dealing with language to make sure that you had not seen **minsocs** before? If you did, you must have conducted the search very quickly, for you have probably heard some 50,000 different words during your life. Yet, somehow, there is the feeling that it is not even necessary to check. You know immediately that these are new words and that this particular string is meaningless in terms of past experiences. But how do you build a memory system that knows so quickly what it does not know?

Or what about this statement?

The minsocs are rakans.

Now you should at least entertain the possibility that what is being said is perhaps meaningful. You still do not know what **minsocs** or **rakans** are, and you know that you do not know. But somehow this statement is treated differently from the first. Perhaps you will hear more about **minsocs** and can add a new word to your vocabulary. Maybe **minsocs** is worth remembering after all.

Notice that the sheer fact that you do not understand the word and cannot interpret it in terms of past experiences does not mean you cannot encode and remember it. If we ask you what **minsocs** are 200 pages from now, you will probably be able to recall that they are **rakans**, even though you do not know the meaning of either of the words. The models we build, then, must be able to decide what things are potentially useful and should be remembered even though it has never seen them before and cannot understand them.

Even fully meaningful statements often make complex demands on memory:

The authors are fascists.

Here all the words are meaningful, yet your reception of the message is hardly a passive, automatic experience. If you take this statement seriously, then to digest it you must retrieve and evaluate enormous amounts of information: What authors are we talking about; what authors would we know that you would also know; what authors do you know who are fascists? Maybe we are talking about ourselves. If so, maybe you should review what you have read so far to decide if there were any statements in the text that suggest that we are fascists. But how do you know a fascist statement when you see one? Whether you can define such statements or not, you will certainly be alerted to them if they occur from here on in the text.

Again there is the question of how to build a memory system that knows just what aspects of the vast amount of data it has stored are pertinent to a particular input. Moreover, it has to be able to find rapidly just the right information needed to evaluate the message, and even if it rejects the message it must be able to use it to guide the interpretation of future events.

Finally, what about the request:

Tell me everything you can think of about authors.

This should open the floodgates. Now you can go on forever, talking about the things you know. You may start talking about authors in general, then move on to specific authors you know, then to their stories, then to the relationships of these stories to your own experiences, and on and on. You probably will not repeat yourself, even though after awhile it would seem that just remembering what has already been said represents an extraordinary feat of memory. We need to build a memory system where all the information is ultimately related to all other information, where new relations linking old concepts can always be found.

We are beginning to see some of the complexity of the requirements for a realistic model of human memory. These illustrations demonstrate the crucial difference between a memory that passively records what it receives and one that actively interprets and analyzes the incoming information. This is precisely the issue that has to be faced in order to develop a convincing model of human memory. The model must be one that can use a conceptual structure to interpret the information it receives, that can compare the incoming messages with what it knows, and that can evaluate the plausibility of something in terms of its past experience. The time has come to see what is involved in actually constructing a model with these fundamental properties of human memory.

A MODEL OF MEMORY

So far we have been talking about general principles, not specific models. Now it is time to attempt to translate this general discussion into a concrete model. We are going to try to build a memory system that has some of the characteristics of human memory.

In building such a model, an important distinction must be kept in mind. There are really two parts to memory, each equally important. One part is the *data base*, the part of the structure where the information in the memory system is actually stored. The data base must be able to encode and remember concepts and events and complex interrelationships—the stuff of human memory. Our first job is to work out the rules of the data base: We do this in the next section.

The other part of memory is the *interpretive process*, the system that uses the information stored in the data base. It is responsible for evaluating inputs to the memory, for storing new information, for answering questions, for retrieving information to solve problems, speak, think, and guide the daily operations of life. The investigation of these interpretive processes will be taken up in the next chapter.

Remembering concepts

Human memory contains an enormous variety of concepts that can be retrieved and used at will. People have concepts of houses, dogs, cars, communists, and Cub Scouts. Most of the time, but not always, labels are attached to the concepts, such as those just used. In addition to the label, large amounts of information associated with any given concept can be produced on demand. The first job, then, is to decide how to represent concepts in a memory system.

Think of a word, say, **teapot.** Ask a friend to explain what it means, or explain it to yourself out loud. What kinds of information do you produce when describing its meaning? A typical explanation looks something like this:

> **Teapot,** *n. A container something like a kettle, made of metal or china.* [*The Golden Book Illustrated Dictionary for Children*]

Or

> **Teapot,** *n. A container with a handle or spout for making or serving tea.* [*The Thorndike-Barnhardt Comprehensive Desk Dictionary*]

Similarly, for other words, say, **tapestry, tart,** and **tavern,** the dictionary states that

> a **tapestry** is *a piece of cloth with figures or pictures woven into it;*

a **tart** is *a small, baked crust or pastry filled with fruit jelly or jam*;

a **tavern** is *a place where beer, wine, or other alcoholic drinks are served.*

These examples remind us that the definition of a word consists of other words. Typically, a definition starts off by saying, "Concept A is really something else, namely concept B"; a tapestry is a piece of cloth, a tart is a pastry. It then goes on to specify the restrictions on the concept. Unlike other places, taverns serve beer and wine. The unique thing about a tapestry is that it has figures or pictures woven into it. A teapot has either a handle or a spout and is used for tea. (The child's dictionary is a bit ambiguous about why teapots are different from other containers: It suggests that the fact that a teapot is made of metal or china is a critical property.)

Another form of information that is often used in explaining a concept is an example. If you were explaining a **tavern** to a friend, you would probably point out some specific examples. If we look up **place** in a dictionary, we might find little else but examples.

Place, *n. A city, town, or area.*

(Notice that the dictionary did not mention **tavern** as an example of a place.) Similarly,

Container, *n. Anything with the property of holding something, such as a box, barrel, can, or jug.*

It seems, then, that if a person or dictionary is questioned about what a word means, all that is produced is other words. For some reason, this does not seem to disturb us. The hoax only becomes apparent if you persist and ask for the definitions of words that are used to define other words. If you go to the *Oxford Dictionary* to find out what a **son** really is, first by looking up a definition of son, then looking up the definitions of the words used to define it, etc., you will trudge through a complicated maze, finally coming to a dead end. You will have gone in a circle, finding **child** defined as an **offspring,** and **offspring** as a **child.**

An important part of the meaning or comprehension of a concept, then, must be embedded in its relationships to other concepts in the memory. On examining the format of typical definitions, a rather small number of relationships seems to predominate: the class of *concepts* to which it belongs, the *properties* which tend to make that concept unique, and *examples* of the concept. A standard definition, then, can be summarized schematically as in Figure 10-3.

FIGURE 10-3

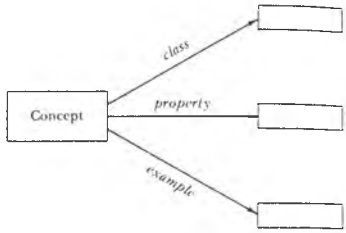

FIGURE 10-4

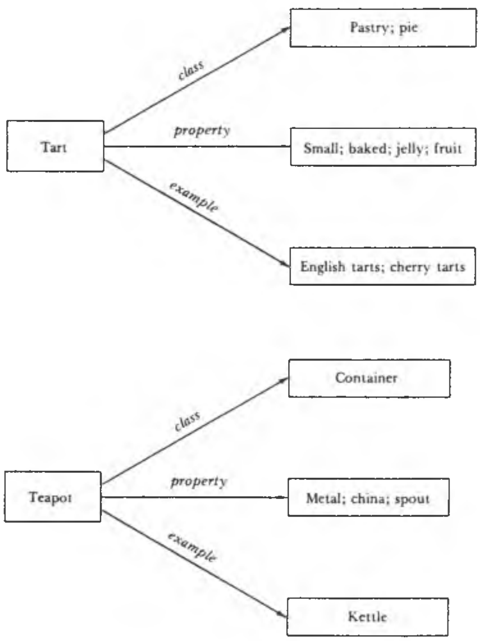

Filling in the blanks with some of the above definitions produces Figure 10-4. Moreover, the words used in the definition ~~are~~ represent themselves concepts, and therefore defined in the same way. The result is an interlinking structure that may not be apparent when looking up definitions, but that certainly becomes obvious when structure is shown graphically, as in Figure 10-5.

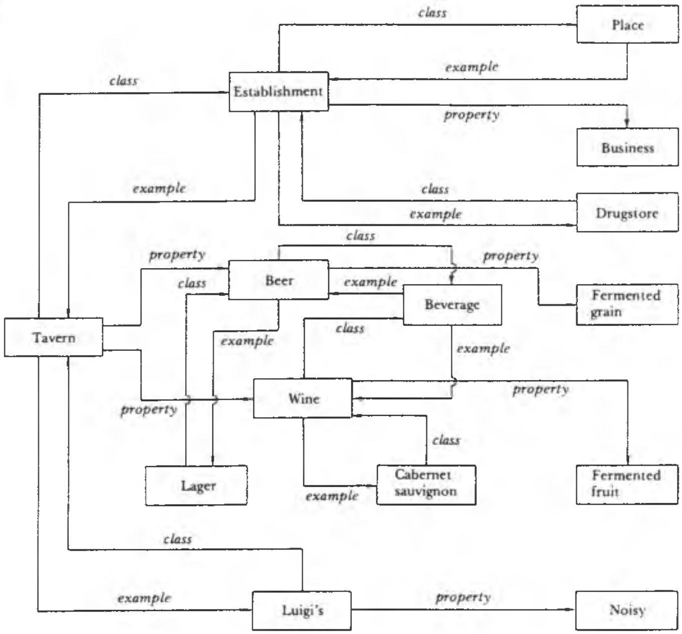

FIGURE 10-5

To represent concepts in the memory the diagrams show two kinds of things: boxes and arrows. The boxes represent the concepts. Notice that the arrows have two important properties. First, they are **directed.** That is, they point in a specific direction. We can follow them in either direction, but they mean different things. Second, they are **named:** there are three kinds of names so far—*property*, *example*, and *class*. Certainly, we need more than these simple relationships if we are to encode things more complicated than concrete nouns in the memory system. The first step in expanding the system is to change the names of the directed arrows to allow almost any action or relation to serve the purpose.

This is easy to do. First, replace the *class* arrow with the verb *isa*. *Isa* is obviously a compound word, constructed of the two individuals *is* and *a*. *Isa* always connects with an object, as in **John** *isa* **man.** Second, replace the *property* arrow with one of two verbs, *has* or *is*. *Has* is used primarily when properties are objects, such as an **animal** *has* **feet.** *Is*, on the other hand, is used primarily when the property is a quality, such as **John** *is* **hungry**, or **Rover** *is* **fat.** (Be careful not to confuse *isa* with *is:* They are quite different.)

Finally, notice that examples are almost always related to class names: The two simply go in opposite directions. Thus, if the class of tavern is an establishment, an example of an establishment will be tavern. So why bother with specific example arrows: Simply let an example be given by following the *class* or *isa* arrow in the backward direction. In summary:

class	*isa*	As in **John** *isa* **man.**
property	*has*	As in **Animal** *has* **feet.**
property	*is*	As in **John** *is* **tall.**
example	*the reverse direction of isa*	As in **John** *isa* **man.**

With these equivalences in mind, the definition of tavern can be redrawn:

FIGURE 10-6

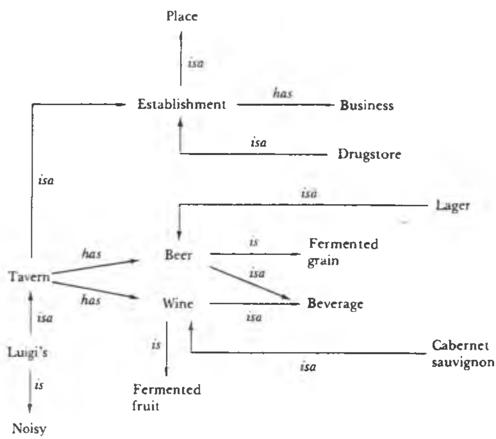

This is a more satisfactory description of the interrelationhips. First, it is much simpler than the original, always a virtue. Second, it captures some of the concepts better. Instead of saying that one of the *properties* of taverns was wine, a peculiar use of the word *property*, we now say simply that a **tavern** *has* **wine**: a neater way of doing things.

In considering the way information is being represented so far, you may have been bothered by the apparent circularity of things. Things get defined in terms of themselves. Moreover, sometimes, things belong to several classes. The memory may indicate that a **dog** *isa* **pet** and also that a **dog** *isa* **mammal**. Which is it? Actually, circularity and apparent lack of precision are desirable in a model of memory, since that is exactly how human memory is. To one person, a pet may be any animal that is domesticated, so that a dog is simply an example of a pet. But another person might have grown up around a household which only had a dog, and that is all he ever saw as an example of an animal that could be domesticated. Hence, when asked what a pet might be, he would reply, "Well, a pet is a dog that is domesticated." Later on, as he grew up, he would broaden his definition to, "A pet is a dog or cat that is domesticated." After many years of this, all the while broadening the definition, he might suddenly have an insight, realizing that, really, a pet is an animal that is domesticated so that, for most purposes, it is a dog that *isa* pet, not the pet that *isa* dog. But he already has the memory structure built up, so he adds the new concept to the old. This is not the neat, systematic logic that language and experience ought to have. But we are describing real behavior, and that is often complicated, confusing, and circular.

FIGURE 10-7

A pet isa dog or cat that is domesticated

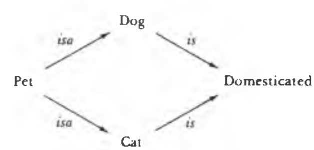

A pet isa dog or cat, and a pet is domesticated

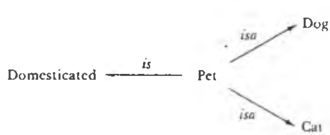

Cats and dogs are pets. A pet isa domesticated animal

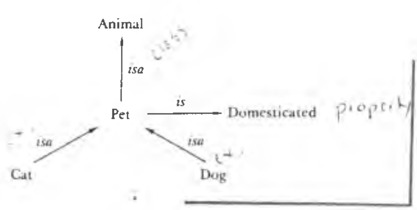

Primary and secondary concepts
needed Later

In considering how to represent various kinds of information in this memory, an important problem comes up. Suppose we are trying to remember the information:

Leo, the hungry lion, has a sore mouth.

The difficulty here comes from the way in which we add the fact that the lion has a sore mouth:

FIGURE 10-8

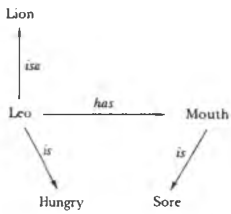

This is one way of representing the sentence. Note that **sore** has to modify **mouth.** If you put the arrow **is sore** off of **Leo,** then all of **Leo** would be sore, not just his mouth.

This description might be all right if this were the only time the concept of **mouth** was mentioned in the memory. But suppose we also know **John is a person who has a big mouth.** Simply adding this information as before would produce

FIGURE 10-9

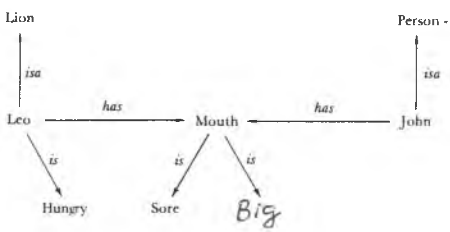

Clearly, this is wrong. When retrieving information about **John,** the memory would think that John's mouth was sore and that **Leo has a big mouth** as well as a **sore mouth.**

The way out of this dilemma is to realize that we need to have only one definition for the concept of mouth but that we need to have many instances where the concept is used, perhaps in modified form. The first definition, the basic one, is called a *primary definition.* The other, the particular use of the concept, is called a *secondary definition.* We represent secondary concepts by enclosing them in angular brackets, like this: ⟨**mouth**⟩. This can be read as "this mouth." This primary–secondary distinction is invaluable, as Figure 10-10 shows.

FIGURE 10-10

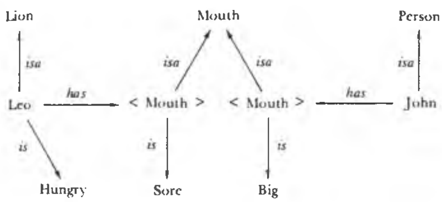

Actually, it is not even necessary to insert a name inside the secondary concept, for its definition can always be found simply by following the *isa* arrow, as illustrated in Figure 10-11. The secondary node labeled

FIGURE 10-11

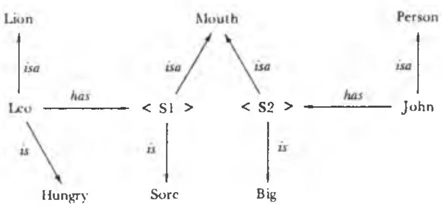

S1 is that of a **mouth**, in this particular instance, a **sore mouth**. Secondary node S2 is also that of a **mouth**, but in this instance it is a **big mouth**. When retrieving information, we can automatically substitute **sore mouth** for S1 and **big mouth** for S2 to recover the correct information in each case.

To introduce the idea of a data base for storing information and its interrelated structure, we have restricted ourselves so far to descriptions of concrete nouns and three basic kinds of relationships: *isa*, *is*, and *has*. These concrete concepts are an important part of the human memory, but they represent only part of the information people normally encounter. What about events? What about the memory for the plot of the last novel you read? How can actions be represented in such a system?

Remembering events

Using the same basic strategy, it is rather easy to add different types of information to the data base. Only two more steps are required, one very simple, the other rather complicated. The easy step is simply to expand upon the allowable types of arrows that can interconnect concepts. Before letting these arrows proliferate freely, however, it is important to decide on the types of arrows that might be connected to events.

The problem is to represent an event in the memory system. We do that by adding a new type of node to the memory, an *event* node. Thus, in the situation

The dog bites the man.

we wish to add the description of that event. To do that, consider an event as a *scenario*, with actions, actors, and stage settings. All the information must get encoded, with each part of the scene properly identified to its role in the event.

Consider again the situation **The dog bites the man.** Here the sentence that describes the event can be broken down into three parts: a subject (**dog**), a verb (**bites**), and a direct object (**man**). But we do not really wish to determine subjects, verbs, and direct objects, for these are often misleading. Take the sentence, **The man is bitten by the dog.** What do you call the subject? **Man** or **dog?** We want it to be **dog.** The instigator of the action is **dog,** so it is **dog,** not **man,** that is our subject.

To record events, some new concepts must be defined. Consider how an event is described. What we want to do is to break it down into a set of simple relationships that describes the basic concepts of the event. Events can often be described in sentences, but the sentence must be analyzed with some care. Linguists are very careful to distinguish among several levels of language. One, called *surface structure,* represents the part that is visible: The actual sentences people speak. The other level is called *deep structure* or *semantic space,* and it represents the meanings that underlie the sentences. Clearly, the important thing for memory is deep structure, or semantic space. Some sentences can look very similar to one another at the surface structure, but mean completely different things at the semantic level. Consider the sentences

Mother is cooking.

Supper is cooking.

These two sentences look very much the same, but they mean quite different things. In one case, Mother is standing at the stove cooking something. In the other, we can hardly imagine supper to be standing in the kitchen cooking something: It is supper that is being cooked, perhaps by Mother:

Mother is cooking supper.

To discover the basic structure of an event without being misled by the surface structure of the sentence that describes it, we always start by ignoring the details of the sentence and by identifying the *action*.[3] The first step in the analysis is to decide what the *scenario* is: What is the action? Next, find the actors and the things being acted upon. The actors, who cause the action to take place, are called *agents*. The things being acted upon are *objects*, and the person who receives the effect of an action can be called a *recipient*. Here are some examples:

Mother is cooking.

Action: cooking

Agent: Mother

Object: none

Frame or template Specifically: case frame for verb "cooking"

Supper is cooking.

Action: cooking

Agent: none

Object: supper.

Mother is cooking supper for Hubert.

Action: cooking

Agent: Mother

Object: supper

Recipient: Hubert

Identifying things this way simplifies life considerably.

Now we see how to represent events in the data base. The entire event centers around some action, so the action becomes the central node: We represent it in diagrams as a node, drawn as a circle around the word (usually a verb) that describes the action. Then the actors and objects that comprise the scenario are attached to the event node by arrows that identify their role: The basic format is shown in Figure 10-12.

[3] These examples and analyses come from Fillmore (1968) in Bach and Harms (1968).

FIGURE 10-12

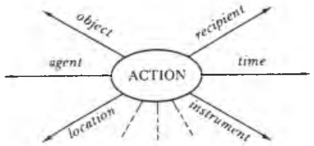

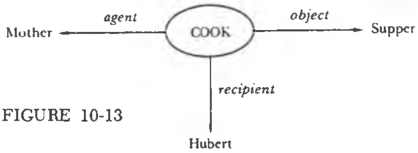

FIGURE 10-13

Thus, the sentences

Mother is cooking supper for Hubert.

and

Hubert's supper is being cooked by mother.

are both diagrammed as the same scenario—that of Figure 10-13. Thus, although the sentences look quite different from each other (they have different surface structure), they have the same meaning (the same deep structure), so they are drawn the same in terms of the information recorded in the memory. Moreover, there is the strong implication that the cooking is being done somewhere (a *location*), with something (an *instrument*), and at some specific time (*time*). These unstated concepts are simply added to the event node whenever they become known. No new structure need be created for them.

Other *cases* (that is what things like *agents, objects,* and *recipients* are called) that are useful are such things as

time: when an event occurs, often specified simply as past, present, or future, or conditional. (*Jack* **kissed** *Louise: Time is* **past**).

location: where an event takes place (*Bob hit Jack* **on the head:** *Location is* **head**).

instrument: the thing involved to cause the event (*Bob hit Jack on the head* **with a rock:** *Instrument is* **rock**).

truth: whether the event was true or false (*I* **did not see** *Jack: Truth is* **not**).

A complete list of the cases used to describe events is given in Table 10-1.

Table 10-1 The Parts of an Event

A more complete frame

Action	The event itself. In a sentence, the action is usually described by a **verb**: The man was **bitten** by the dog.
Agent	The actor who has caused the action to take place: The man was bitten by the **dog.**
Conditional	A logical condition that exists between two events: A shark is dangerous **only if** it is hungry. John flunked the test **because** he always sleeps in lectures.
Instrument	The thing or device that caused or implemented the event: The **wind** demolished the house.
Location	The place where the event takes place. Often two different locations are involved, one at the start of the event and one at the conclusion. These are identified as **from** and **to** locations: They hitchhiked **from La Jolla to Del Mar.** **From the University,** they hitchhiked **to the beach.**
Object	The thing that is affected by the action: The wind demolished the **house.**
Purpose	Identifies the purpose of the event: Jack took Henry to the bar to **get him drunk.**
Quality	A descriptor, one that modifies a concept: The surf was **heavy.** There were **93** people in class.
Recipient	The person who is the receiver of the effect of the action: The crazy professor threw the blackboard at **Peter.**
Time	When an event takes place: The surf was up **yesterday.**
Truth	Used primarily for false statements: I do **not** like you, Hubert.

The event

Yesterday, at the beach with my new camera, I photographed the house on Ninth Street.

is analyzed as

action: photograph

agent: I

object: house on Ninth Street

location: beach

instrument: my new camera

time: yesterday

This analysis can be broken down even further. The *object* can be analyzed as a *concept* (house) plus a *location* (Ninth Street). The *instrument* is a specific camera, namely mine. Hence, the final structure is as shown in Figure 10-14.

FIGURE 10-14

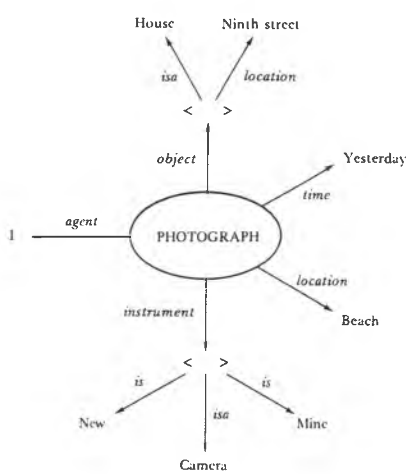

All these relationships are entered into the data base much the same way that concepts were, except that now there is a richer set of possibilities than simply *isa, is,* and *has*. There is one simplification of the structure that is sometimes useful, *shorthand notation*. Often in the case of a simple event, such as **Mother cooks supper,** there is little ambiguity as to the role played by each concept (**Mother** and **supper**). In this case the full event notation, that shown in Figure 10-15*a*, is not needed. Rather, the simplified structure of Figure 10-15*b* can be

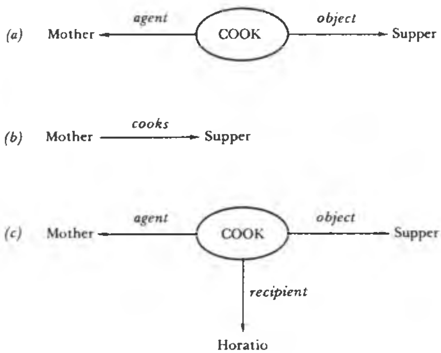

FIGURE 10-15

used. Notice that this shorthand notation simply uses the action as the name of the relation (arrow) that connects the two concepts. There is no difficulty with this shorthand as long as it is recognized that the two versions shown in parts (*a*) and (*b*) of the figure are equivalent. As soon as a new detail is added to the description of the event, however, then the shorthand notation no longer works. Hence, **Mother cooks supper for Horatio** must be described in the full notation of Figure 10-15*c*.

Now that the full power of the data base is starting to emerge, add these events to the data base described earlier in Figure 10-6.

Bob drinks beer.

Mary hit Louise hard yesterday at Luigi's.

Al owns Luigi's.

Bob likes Louise.

Al's dog, Henry, bit Sam because he yelled at Mary.

Louise drinks wine.

Mary likes Bob.

When this is done, the data base is considerably enriched, since now not only are many concepts defined, but we can begin to see the events that take place involving those concepts. A sample of how the final data base might look is shown in Figure 10-16. Note that here both

FIGURE 10-16

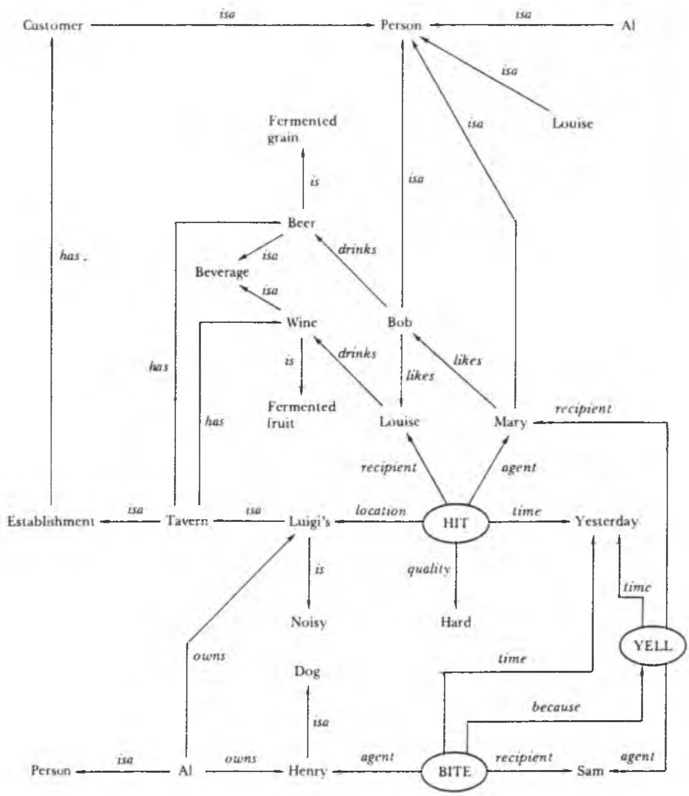

shorthand and full notation are used to describe events, but there should be no confusions caused by the combination. (A few new techniques are illustrated in this figure, so it would pay to examine it with care.)

We now have the basic design for the data base underlying human memory. The memory system is an organized collection of pathways that specify possible routes through the data base. Retrieving information from such a memory is going to be like running a maze. Starting off at a given node, there are many possible options available about the possible pathways to follow. Taking one of these paths leads to a series of crossroads, each going off to a different concept. Each new crossroads is like a brand new maze, with a new set of choice points and a new set of pathways to follow. In principle, it is possible to start at any point in the data base and, by taking the right sequence of turns through successive mazes, end up at any other point. Thus, in the memory system all information is interconnected.

As it stands now, the memory is passive. It is a network of potential pathways, each of which could, in principle, be used by interpretive and retrieval processes. It is now time to examine some procedures for retrieving and manipulating the stored information, the procedures that describe how the memory is used.

As an introduction to the problem, answer the following queries from the data base shown in Figure 10-16.

Query: **Do people drink beverages?**
Query: **Does Al like Mary?**
Query: **Is Louise a customer?**

SUGGESTED READINGS

Suggested readings for Chapters 10 and 11 are combined at the end of Chapter 11.

11

Memory processes

MECHANISMS FOR INPUT AND OUTPUT
Processing an input
Making responses
DEDUCING A PERSON'S MEMORY STRUCTURE

SEARCHING MEMORY
Multistage search processes
Subset relations
Overlapping relations
Disjoint relations

REORGANIZING THOUGHTS
Correcting errors
Evaluating an input
VIEWING THE DATA BASE

UNDERSTANDING VERSUS REMEMBERING
MEMORY AS INFORMATION PROCESSING
SUGGESTED READINGS

In the preceding pages, we developed a basic organizational structure for memory, including methods for encoding concepts and events. The structure of memory is only half the story. It is time now to explore the kinds of cognitive processes that might operate on such a memory structure.

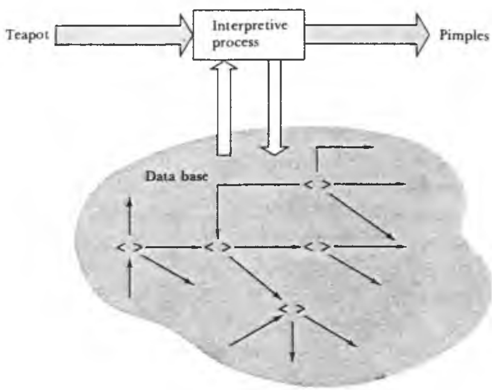

FIGURE 11-1

In this chapter, we explore the second half of the memory process: the *interpretive process.* The story is only a guide, since it soon gets entangled in the complexities of human thought. The general approach may be on the right track, but the details need careful scrutiny, both in terms of the logical properties of the processes involved and the experimental tests of their implications. The study of the cognitive processes of memory is new: The principles we discuss can be no more than a mere beginning.

MECHANISMS FOR INPUT AND OUTPUT

A memory system must communicate with the world. It has to be able to take in statements and recode them into a format suitable for storing in the memory. It must also be able to respond to questions. Given the kind of memory structure described so far, the basic strategy for handling inputs and making responses is relatively straightforward.

First let us consider two simple versions of the problems. The first problem of input is the problem of adding new statements into the data base; the first problem of output is to translate the information from the memory system into a coherent set of statements that describe the data that have been stored.

Processing an input

Suppose that Mary has just walked into our lives. We perceive her as a simple concept:

Mary is a fat, pimply person.

This example is not diagrammed

To get the concept into the data base requires that it first be translated into the standard format by which information is represented. That is, the concept of Mary as we now know her resides in three simple statements:

Mary isa person.

Mary is fat.

Mary has pimples.

At this point, we are concerned with the mechanisms that add these three statements to the data base. We assume that the concept of Mary, as perceived by the various perceptual and pattern-recognition processes, results in an image of Mary that contains these three statements. Now, the task is to translate the image into the data base. The illustrations

FIGURE 11-2

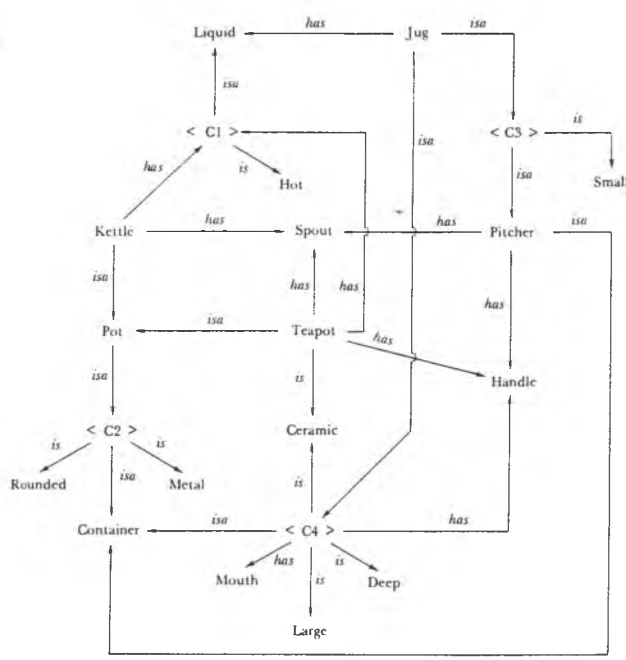

we use are all linguistic—Mary is described in simple sentences—but that is for convenience only. Presumably the actual memory of **Mary** as **fat** and **pimply** is nonverbal. The verbal description of the data base is the easiest to illustrate, but it is important to realize that these verbal descriptions are only representations of the actual, symbolic processes that must actually occur.

Making responses

Let us start with the simplest situation: a set of concepts. The data base shown in Figure 11-2 contains the definitions of various kinds of pots and containers. Notice that all the concepts do not have names. In particular, there are intermediate concepts, labeled C1, C2, C3, and C4 in the figures, that are used during the definition of other entries. These intermediate concepts are the **secondary** concepts discussed earlier. If you follow the definitions through, you should have no difficulty in understanding their use.

How does this system answer questions about what it knows? First, try some simple ones.

Query: **Tell me about a teapot.**

The memory system should respond something as follows.

> A teapot is a pot. It is ceramic. It has a handle, a spout, and a, let me see, a liquid that is hot. A pot is a, um, a container that is rounded and metal. A kettle is a pot. Now let me tell you about a kettle. A kettle has a spout and, let me see, a liquid that is hot. Now let me tell you about. . . .[1]

This output is simple. It represents a straightforward attempt to follow the paths leading from each node. No interpretation is involved. But what happens when simple reasoning is required?

Using Figure 11-2, answer the following queries of the memory system.

A. *Query:* **Does a jug have a handle?**[2]
B. *Query:* **Is a pot made of metal?**
C. *Query:* **Is a teapot a pot?**
D. *Query:* **Is a teapot ceramic?**
E. *Query:* **Does the answer to D contradict the inference resulting from the answers to B and C?**

[1] In this hypothetical illustration we indicate something of a possible strategy for recall. First, all the information around a node is described, then one of the nodes referred to is described, and so on. Whenever a second node is reached, however, it must be examined with some care, first following the *isa* arrow to find its name, and then following the others to add qualifiers. This requires time and effort, so we assume the "person" fills in the time by such innocuous phrases as "let me see," "um," and ". . . you know."

[2] The correct answer is yes.

DEDUCING A PERSON'S MEMORY STRUCTURE

Within the memory system, everything is eventually related to everything else. Thus, if retrieval were allowed to continue, it could go on forever. But the structure of the output is intimately related to the structure of the data base. This fact can be used in the clinical assessment of patients. Several possible techniques can be used. In one, the technique of *free association,*

FIGURE 11-3

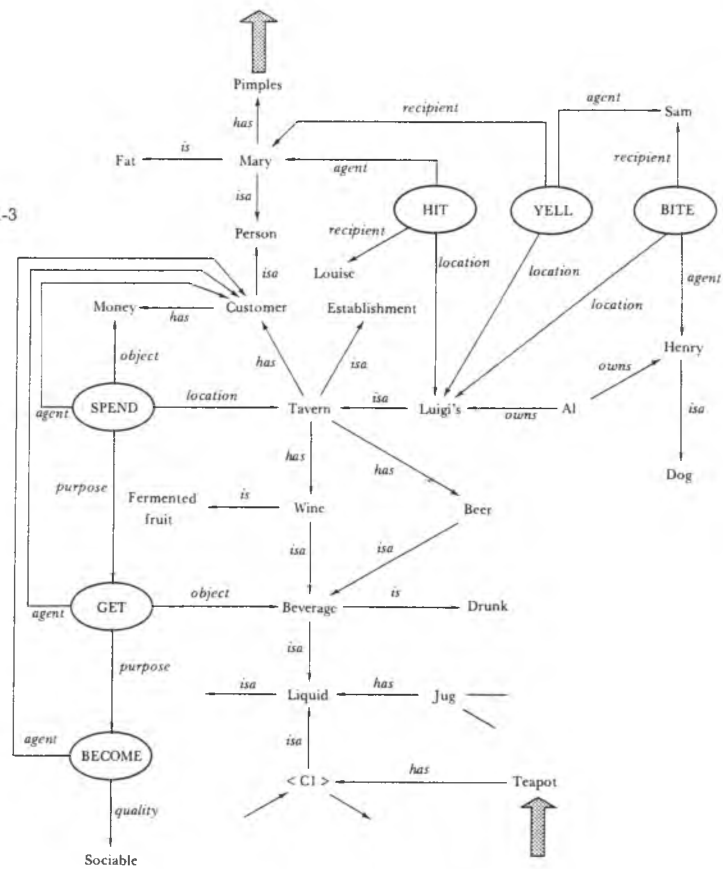

the patient is presented with a word and asked to respond with "the first word that comes to mind." In variations of this technique, the patient may be asked to discuss any topic that he wishes, whether it seems relevant or irrelevant to the problem confronting him. With the interlocking structure of the data base, everything must have some relevance to the consideration of everything else.

Look what would happen if the more complete version of the data base of Figure 11-3 is used to produce free associations and discussion of the word **teapot.**

Start with **teapot.** The path leads to **liquids** and **beverages.** Now the path goes two different ways.

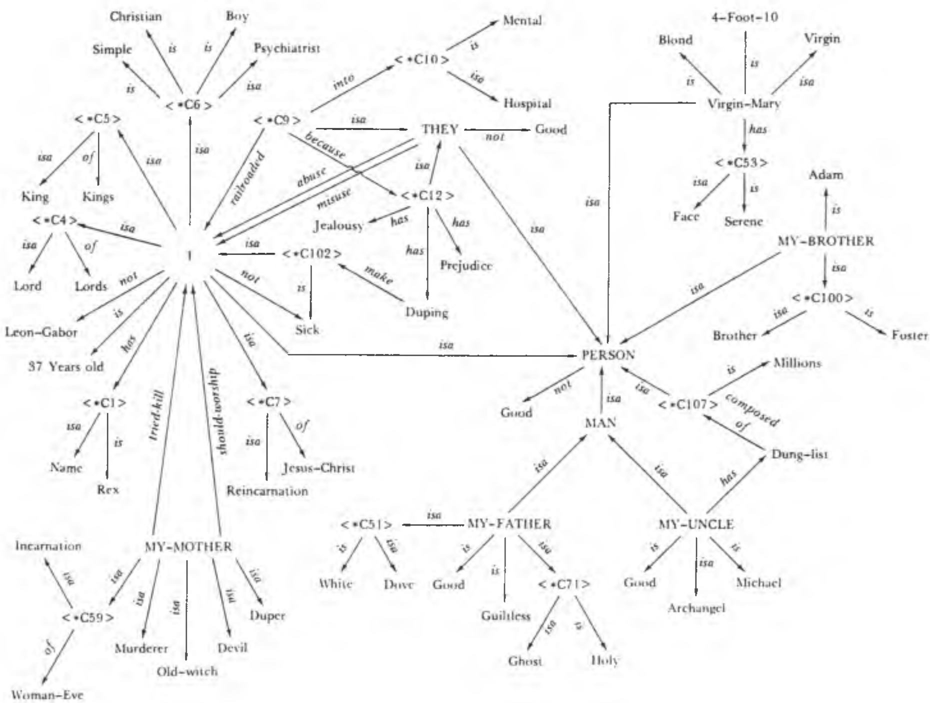

FIGURE 11-4

> *Customers* spend money at taverns to get *beverages* for sociable purposes.

Alternatively, the path reads

> **Wine and beer are** *beverages*. **The tavern, Luigi's, has wine and beer, as well as** *customers*.

Now, both paths have gotten us to **Customer**. **Customer** takes us easily to **person** and, hence, to **Mary** (the girl with **pimples**). But we are very close to a whole, interlocked series of events, the incident at **Luigi's**.

The type of response that this data base produces is determined to a large extent by the interpretive system. To whom does the data base belong? Is it the data base for **Sam**? If so, presumably he is very much concerned with that incident. With **teapot** as a starting point, he easily remembers either **customer** or **tavern** and then **Luigi's**. But once at **Luigi's**, he is attracted by the series of events in which **Henry** bit him, evidently because he yelled at **Mary**. But **Mary**, that **fat girl** with **pimples** had hit **Louise**. (**Louise**, we will discover later, is **tall** and **handsome**.)

Suppose **Sam** were asked to free-associate to the word **teapot**. Is it not possible that he responds **pimples**? The skilled clinician tries to make use of that bizarre association to probe the underlying memory structure, hoping thereby to deduce the complete picture. In this example, the structure is a rational one. With many mental ailments, of course, things are not put together so rationally. You might enjoy probing this structure of Leon-Gabor, a patient who believed himself to be Jesus Christ.[3]

Figure 11-4

SEARCHING MEMORY

Query: **Are all alps mountains?**
Query: **Does a canary have blood?**

One way of testing memory is to see how people answer queries. Actually, the most valuable thing to observe is how quickly they can give answers, not what answers they give. To see this, let us examine another segment of the data base. This one is copied from the work of Allan Collins and M. Ross Quillian (1969) except that is has been redrawn to use the notation described in this book.

Consider the following queries:

1. **Is a canary yellow?**
2. **Does a canary have wings?**
3. **Does a canary breathe?**

How would you answer these by using Figure 11-5? Begin with the most difficult question: 3. **Does a canary breathe?** The obvious first

[3] The structure for Leon-Gabor was extracted from the case study by Rokeach, *The Three Christs of Ypsilanti* (1964).

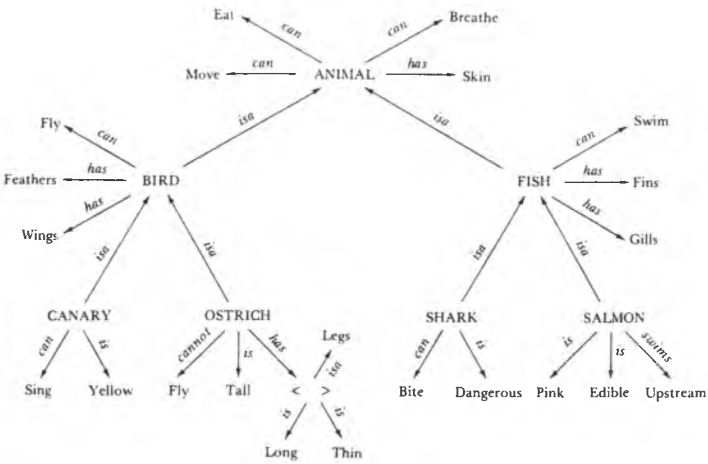

FIGURE 11-5

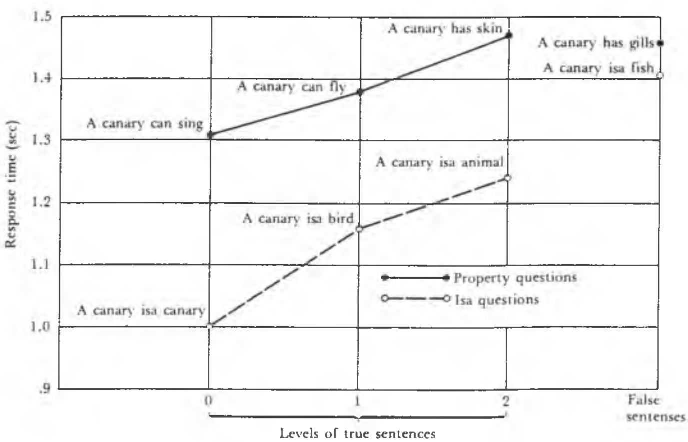

FIGURE 11-6

From Collins and Quillian (1969).

step is to look up the properties specifically associated with a canary and see whether or not breathing is part of the description. If breathing is found on the list of properties, you can immediately respond yes. But if breathing is not there, you cannot conclude that the answer is no. It is quite possible that **breathing** is stored with bird or animal: After all, why should not common properties modify the entire class to which they belong?

What *isa* canary? **A canary isa bird.** The next step, then, is to check the definition of birds: **Does a bird breathe?** Again, there is no information about breathing, so the upward search must be continued. **A bird isa animal.** What do we know about animals? Well, an **animal can breathe.** Hence a **bird** (which *isa* animal) **can breathe;** hence a **canary** (which *isa* bird) **can breathe.** If this is the way information is represented in the memory, then it should take much longer to find the answer to a query like **Does a canary breathe?** than to one like **Does a canary have wings?**

This is what is tested: the time taken to answer these types of questions. Subjects are asked to respond to a number of sets of such queries and the times required to make a response is measured. The results are shown in Figure 11-6. Here, several types of sentences are shown. One sentence, **Isa canary a canary?** is included for calibration purposes. It shows how long it takes someone just to read the sentences and make an answer, even when no real search of memory is required. In addition, two different types of sentences are distinguished. One type, called an *isa* question, needs only to follow the connection labeled *isa*. Hence, **isa canary a bird?, isa canary an animal?, isa oak a tree?** The other type, called property questions, checks properties such as **has skin, can breathe, has acorns.**

The data shown in Figure 11-6 indicate that the more concepts needed to be checked, the longer it takes to get an answer. In fact, it would appear to take around .1 sec for each extra level in the memory that must be searched. Note that it takes about 1.0 sec just to answer **yes** to the statement **a canary isa canary.** This is the time needed to read the sentence, make a simple decision, and push the response button. A sentence that requires searching one level up the *isa* chain (like **a canary isa bird**) takes about 1.18 sec, so we can conclude that the extra .18 sec is due to memory search time. A sentence that requires the memory search to follow two levels of *isa* statements (like **a canary isa animal**) takes about 1.24 sec, or an additional .24 sec over that time required by the most simple sentences.

Similarly, for the case where property lists must be searched, the most direct search—**a canary can sing**—takes about 1.3 sec, with an extra .8 sec required to discover that **a canary has skin.**

These experimental results are only preliminary: They cannot yet be used to do more than give a hint of the structure of memory and how the retrieval process operates. But they do offer some tantalizing evidence of the general nature of the system.

Exercise: Note that two more data points are shown in Figure 11-6—the time taken to respond to false statements such as, **a canary has gills,** and **a canary isa fish.** Now look at Figure 11-5 and try to determine what decision process might be used to answer such questions. How would you deal with such statements as **a canary is dangerous,** or even **a canary is peaceful,** statements for which no information at all is stored, at least not directly?

go to p. 414

Finally, consider the type of data base that would allow you to deal with sentences of the form:

Madrid is Mexican.

A pecan has a castle.

Chicago has mountains.

An igloo would melt in Texas.

Bicycles defeat smog.

Multistage search processes

Query: **All alps are mountains.**

To decide whether this is true or false requires a different type of memory search from those discussed for queries of the form **a canary has skin.** In the canary question, the basic problem was to decide whether the definition structure of canaries was consistent with the statement. With the alps question, we need to examine every single instance of alps: quite a different operation.

A large set of questions of this form has been investigated by David Meyer (1970). He considered subjects' responses to questions of the following four types:

All P are S.

Some P are S.

All S are P.

Some S are P.

Depending upon the relationship between S and P, these questions pose different kinds of search strategies. Consider these cases:

Subset relations. Let **P** be an example of S. Thus, **P** might be an alp and S a **mountain,** or **P** might be a **ruby** and S a **gem.** Hence

All rubies are gems. *True.*

Some rubies are gems. *True.*

All gems are rubies. *False.*

Some gems are rubies. *True.*

Overlapping relations. Let **P** and S be two things that overlap in meaning but neither is an example of the other. Thus, **P** might be **women** and S **writers,** or **P** might be **conservatives** and S **students.** Hence, we get the four types of sentences:

All women are writers. *False.*

Some women are writers. *True.*

All writers are women. *False.*

Some writers are women. *True.*

Disjoint relations. Finally, let **P** and S be things that have nothing in common: They are *disjoint.* Thus, **P** might be **house** and S **vacuums,** or **P** might be **books** and S **cats.** All four versions of these sentences are false.

All books are cats. *False.*

Some books are cats. *False.*

All cats are books. *False.*

Some cats are books. *False.*

Given the kind of memory structure we have been studying, different search procedures are needed to answer these queries. For example, to answer the query:

Are all chairs furniture?

one could start with chairs, find all examples, and check each to make sure it was indeed a type of furniture. If this scheme is followed, the time required to answer the question

Are all thrones furniture?

ought to be a great deal shorter than for the first question. After all, there are far less examples of thrones than of chairs in most people's experiences, so the time needed to do the search should be much less.

Following this reasoning, it is possible to test various theories of the way memory is searched by changing the number of examples that must be searched to answer the questions and seeing how the times taken by subjects to reach their decisions change.

The results from such experiments indicate that there are at least two stages to the process of answering this type of query. First, there is a check to see whether the two things being compared have anything in common with each other—that is, do they intersect? For example, in these two queries:

All wheats are typhoons. (**All S are P.**)

Some chairs are people. (**Some S are P.**)

there is no common relation between S and **P** (**wheats–typhoons**, or **chairs–people**), so we know immediately that the correct answer to these queries is **false.** But if there is something in common between the two concepts, for example:

All pilots are men. (**All S are P.**)

then the first stage of analysis does not produce an answer. In this case, a second stage is needed to analyze whether, without exception, every **pilot** is indeed a **man.** The additional processing, of course, requires additional time, and subjects are correspondingly slower in making a response. Notice that with a question of the type:

Some pilots are men. (**Some S are P.**)

a single positive instance is sufficient for us to answer **true.** Hence the response is fast, since it is given by the first stage of analysis. In general, then, the search for answers to questions of this particular type can be represented as a two-stage search process operating on an organized memory structure. Figure 11-7

The common feature of all our search procedures is that both the retrieval information and deductive logic is required to answer queries of the memory. The answer is seldom stored directly in the memory system. It must be ferreted out, dragged from the corners where it may be hiding, and painstakingly put together. Even so simple a question as

What is a quigee?

requires an extensive search and logical construction of the information found. Figure 11-8

There is a long history in psychology of studies which have demonstrated that memory is seldom a simple storehouse of events, that it is a collec-

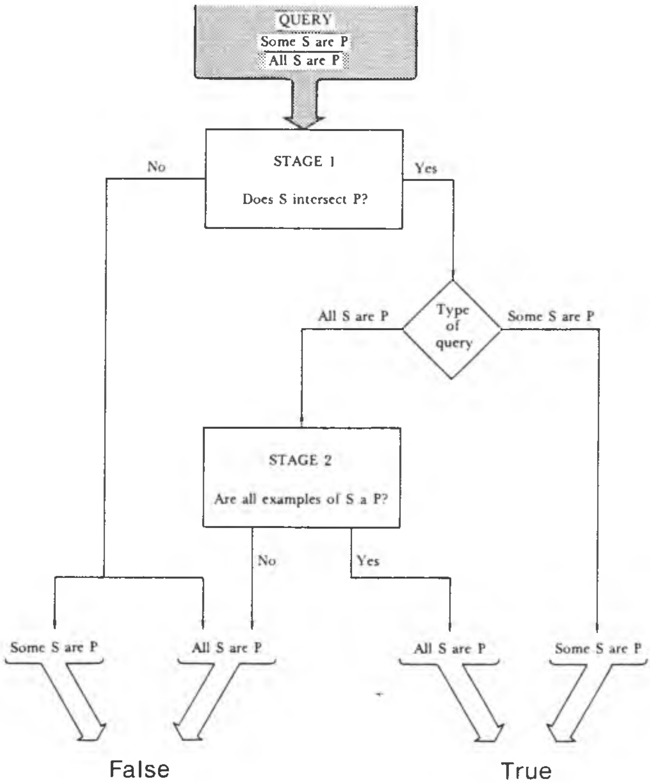

FIGURE 11-7

After Meyer (1970).

tion of ideas that must be worked on to reconstruct the image that is being retrieved. But it is only when we make a careful study of the nature of structure and the retrieval process that might be operating that the truly constructive nature of human memory emerges.

REORGANIZING THOUGHTS

The data we have just been considering indicate that there is a reasonable amount of general structure to the information in the data base. It seems that information is stored where it is most efficient. Rather than have a statement like **has wings** stored with the concept of each and every single bird in the data base, it would appear that the information is stored only once, at the more general concept of birds. This

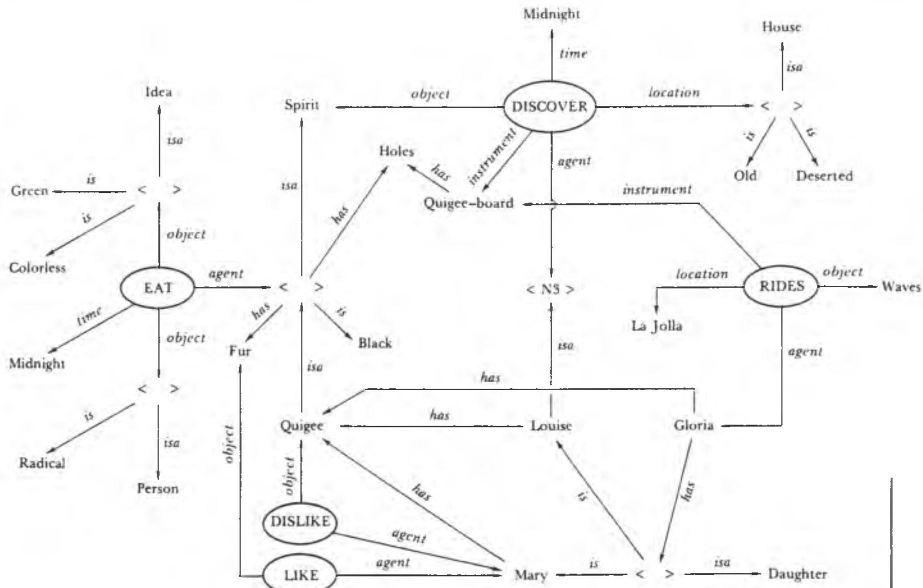

FIGURE 11-8

assumes quite a bit of intelligence on the part of the system, since at times, information within the memory will have to get reorganized. Let us now examine some of the processes by which this reorganization and generalization of information might occur.

Here is a picture of the concepts that might be encoded within the data base of someone's memory. To simplify the discussion, no events are portrayed, simply concepts related to **Mary**, **Bob**, and **Sam**. Given the kind of information shown—that Mary is **short**, **fat**, and **female**, that **Bob** is **red-haired** and **radical**, and Sam is **smart**—what kind of conclusions can be drawn by thinking about the relationships that exist? After all, as more and more information accumulates about concepts, it is perfectly reasonable to stop now and then to ask what has really been learned.

Figure 11-9

Query: **Tell me about person.**

For this query, the memory system should respond with a list of people, and then the properties of those people. But in the process,

FIGURE 11-9

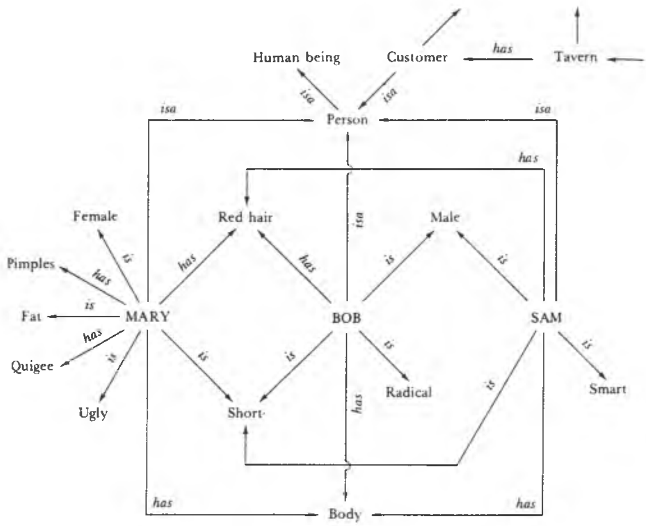

it might learn some interesting things. For example, consider this hypothetical output from the interpretive processes:

> Mary, Bob, and Sam isa person. Mary is female, fat, short, and ugly. Mary has red hair. Hmm, Bob has red hair. Sam has red hair. All persons have red hair. Bob is short. But look at this. Bob, Mary, and Sam are short. All persons are short. . . .

From the information stored in this data base, **all persons** are short, red-haired, and have a body. Thus, the concept of **person** has been generalized by putting together common information.

The system for generalizing is rather simple. First, examine all instances of a concept for information held in common. Whenever the same information is found stored at all concepts, generalize the knowledge of these concepts. Do all people have bodies? The data base only has three examples of people, and each of them does indeed have a body. The obvious thing to do is to remove the property of body from each of those three people and put it in a common location, as the information **person has body.** When we do this, it simplifies the overall structure of the data base.

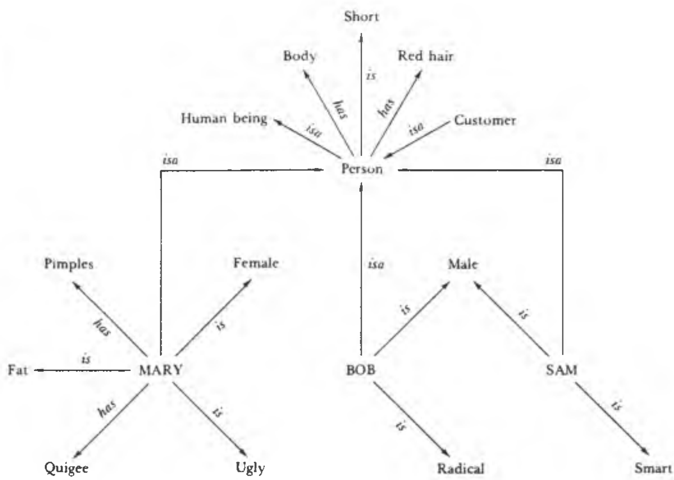

FIGURE 11-10

These generalizations of Figure 11-10 are peculiar: All people have red hair, are short, and have bodies. We would agree with the last, but not the first two. But, in part, this is simply because the memory that is illustrated only knows three people. You disagree that all people have red hair, but this is because you have encountered hundreds or thousands of people, and many of them did not have red hair. For the system illustrated, having red hair and being short is just as accurate a characterization of people as is the fact that they all have a body.

Some further refinement of the memory data-base is possible. Look at Figure 11-10. Here we see that Bob and Sam are persons that are male. This information can be combined to form a new concept, namely, that of N1, a **male person** (see Figure 11-11). For the moment, the memory system only knows that it can find a common concept between some of the persons in its data base, namely the property of being a **male**. Later on, it might learn that it should call the concept **man**, but at the moment, the generalization is indicated with an unnamed secondary node, labeled N1 for convenience.

We now see how to form general concepts based on the analysis of information learned from past experience. The generalization scheme is deceptively simple. As it is applied over and over again, the memory is continuously being reorganized. Properties that concepts have in common tend to migrate upward and become attached to more general items. In time, only the unique or important characteristics of a concept will remain attached specifically to the concept. When the memory system thinks about Mary, it is going to think first about what makes her unique. More general properties can be retrieved by moving upward through the arrows labeled *isa*. With such a generalizing mechanism, the response to a question about a concept might be:

Query: **Tell me about Mary.**

Well let me see. Mary is a fat ugly female with pimples and a quigee. And she is a person, of course, and therefore is short and has red hair and a body. Would you like me to compare Mary with the other people I know?

Correcting errors

Obviously, when making generalizations, the memory system can only use the information it has available at any given instant of time. This is bound to lead to errors, since it may not know very much. This aspect of its behavior is compatible with human performance. But there should be some way of counterbalancing these tendencies to over generalize. A mechanism is needed for redoing concepts as more information comes in. Right now, the memory system has the following information about the concept of person:

go to p. 428

FIGURE 11-11

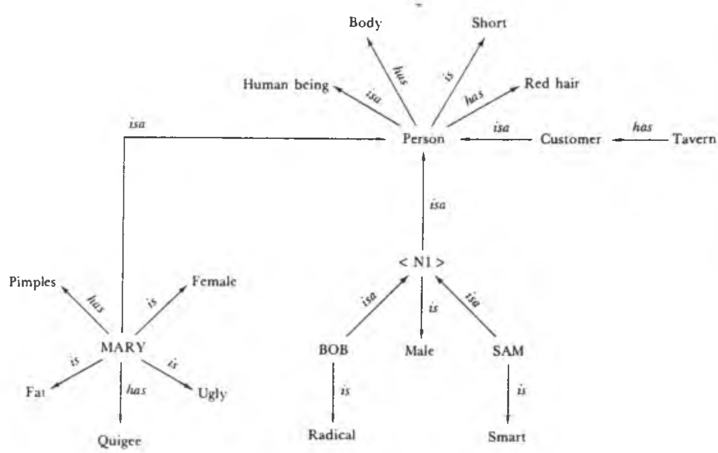

person
  *** isa human being.**
  *** is short.**
  *** has body.**
  *** has red hair.**
  Mary isa *.
  customer isa *.
  N1 isa *.

To simplify the description of the node we let the symbol * stand for the concept that is being defined. Thus *** is a human being** represents the information that a person *isa* human being and **Mary isa *** represents the information that Mary *isa* person. Suppose that the following facts become available.

Person is happy.

Louise isa person.

Harry isa person.

So far, no problem. We add these three facts to the list of things known about person and also start three new concepts for **happy, Louise,** and **Harry.** (At this point, it would be wise for you to get some paper and copy the network in Figure 11-11, adding the three statements above, so that you can modify it as we go along. We are going to go through some reorganizations that may be difficult to follow unless you write them out.)

Suppose we now learn that

Sam is radical.

Sam is ugly.

From this information, a new rationalization can be formed. First add the new information that **Sam is ugly** and **radical** to the definition for **Sam.** Similarly, add **Sam** to the definition for each of these concepts. Note that Mary is also **ugly** but there is nothing much else in common between Sam and **Mary.** However, both Sam and **Bob** are radicals. Moreover, they are both **N1.** Thus, the new generalization—**N1 is radical.**

The memory structure now looks like this:

Sam
  *** isa N1.**
  *** is smart.**
  *** is ugly.**

Bob
  * **isa N1.**

N1
  * **isa person.**
  * **is male.**
  * **is radical.**
  Sam isa *.
  Bob isa *.

FIGURE 11-12

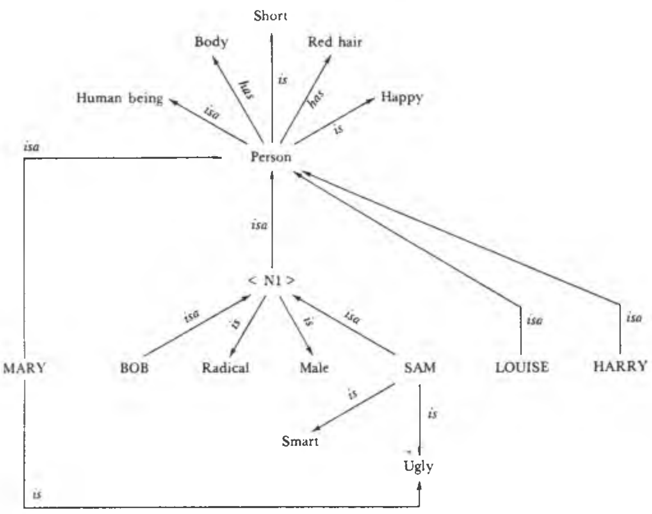

Suppose the memory system continues to learn about people, in particular about **Louise.** So far, all it knows about Louise is that **Louise isa person.** Then the following information is learned.

Louise
  * **isa person.**
  * **is old.**
  * **is handsome.**
  * **is stupid.**
  * **is tall.**

What have we here? **Louise is tall,** but **Louise isa person** and a **person is short.** At this point, the memory system must do something to correct the conflict.

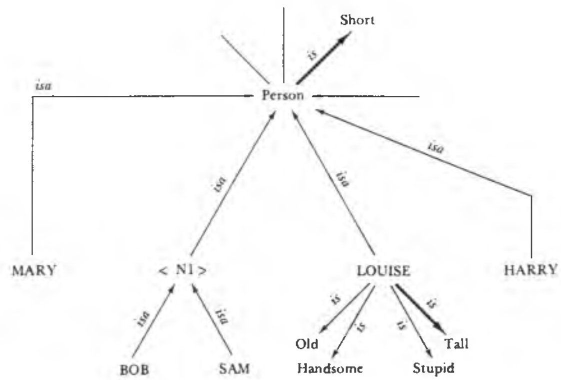

FIGURE 11-13

The first thing it ought to do is make sure it got the proper information and ask its informant:

> Did you really mean to say that **Louise is tall**?

If the information is verified, then it must have some way of correcting the previous overgeneralization. There are several possible strategies. One strategy is to remove the offending information from the higher concept and replace it on all the lower members. That is, delete the statement **person is short** from the definition of **person** and, unless there is conflicting evidence, add it to all the other people the memory knows about, namely **Mary, N1,** and **Harry**. Then, later on, it might reanalyze those individuals to see if it can combine features again.

A second strategy—and the one we follow—is to divide the concept of **person** into two different groups, on the basis of whether they are **short** or **tall**. Thus, we form the concepts N2 and N3:

N2
  * isa person.
  * is short.
  Mary isa *.
  N1 isa *.

N3
  * isa person.
  * is tall.
  Louise isa *.

Now we make the appropriate changes in *person*, ending up with

person
  *** isa human being.**
  *** has body.**
  customer isa *.
  N2 isa *.
  N3 isa *.

What about Harry? All we know for sure is that he **isa person** (and, therefore, like all other persons is a happy, red-haired human being). But suppose we now learn that

Harry is old.

Harry is handsome.

Harry is tall.

Harry is most like **Louise. Harry isa N3**, one of the **old, tall, handsome people.**

FIGURE 11-14

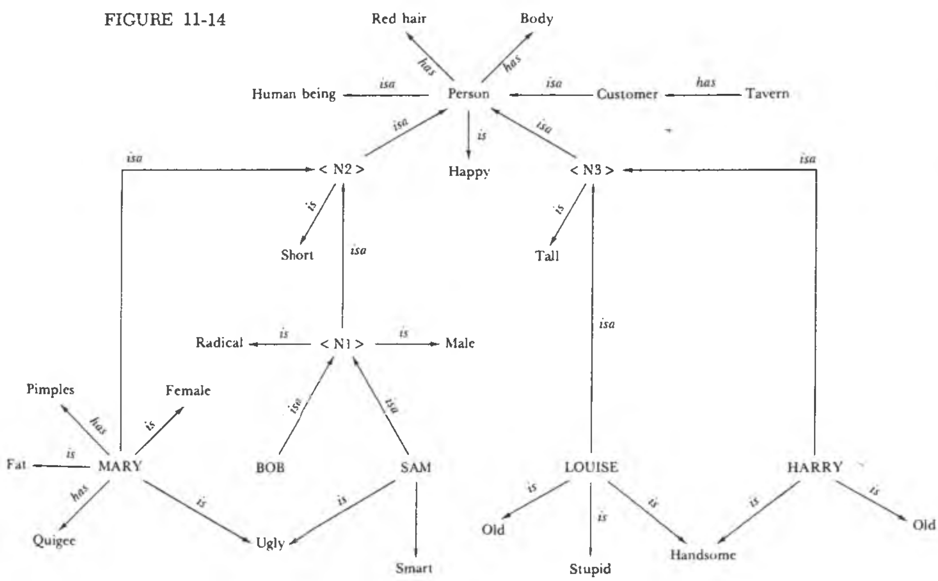

We have now accumulated a rather full list of properties. Here is the new network, in Figure 11-14. We have not yet shown the generalization of information that *Louise* and *Harry* are both *handsome* and *old*. If this were a human memory, the human might very well ponder what it knows from its data base (Figure 11-14) and make the following statements:

> Is Mary a radical? If so, then radicals can be either male or female and there is a good chance that Bob is ugly. Come to think of it, I bet that Harry is stupid. Do radical males have pimples?
>
> I have been thinking about the concept of a person. Anyone who is a person seems to have red hair, a body, and is happy. He is also a human being. There seem to be two major types of persons: One type is tall and handsome and old, the other is short. A short person can also be one of two types, one is female, fat, and ugly, with pimples and a quigee. The other kind of short person is a male radical. Some radicals are smart, but ugly. Some tall, handsome persons seem to be stupid.

Evaluating an input

We are making good progress. In fact, the model is developed far enough to begin considering some of the questions raised in the introduction. How would this memory system carry out some of the retrieval tasks discussed earlier? Remember the statement:

The authors are fascists.

We proposed that you did not passively accept this statement, but rather actively tried to interpret it. We wondered what rules you used to evaluate the input. Now we are in a position to examine that type of evaluation.

Here is the memory with which we will be working.

Figure 11-15

Note that since our last encounter, it has done some reorganization. Moreover, apparently it has received more information about **Mary**, **Bob** and **Sam**, since it has generalized its concepts of **N2**. It has also heard the word **fascists** but knows only that a **fascist is cruel**. Now it receives the statement

A person isa fascist.

This time it is not going to accept the input at face value as before. Now it will try to integrate the incoming message with what it knows.

Informant: A person is a fascist.
Memory: I do not know very much about fascists, but I did not realize that a person in general might be a fascist. Maybe you are

FIGURE 11-15

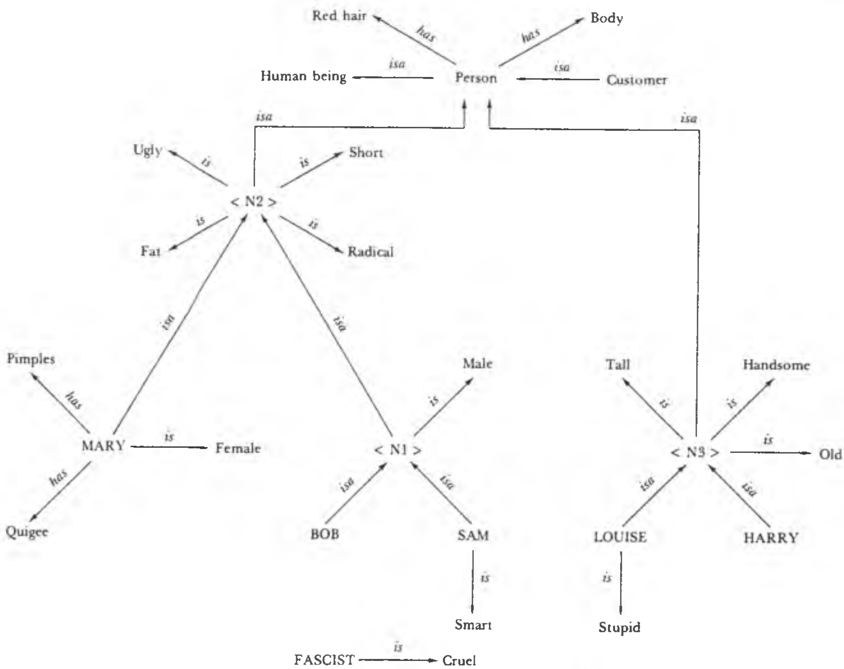

talking about a particular type of person—let me see. Is a fascist tall or short?

Informant: As a matter of fact, a fascist is usually short like Hitler Mussolini, and Napoleon.

Memory: Ah-hah. Is it ugly?

Informant: I would say so.

Memory: And radical?

Informant: In a sense.

Memory: Right. I know some persons who are fascist. I just never connected the name before. By the way, those persons you mentioned before, Hitler, Mussolini, and Napoleon—are they male or female?

Informant: Male.

Memory: Thanks. I will remember that.

The part of the memory that is affected by this interaction is shown in Figure 11-16. The first thing to notice is that an N2 has now been named: N2 is the same as **fascist**. As a result, by implication, the data base now encodes that **fascists** tend to be **fat**, **short**, **ugly**, and **cruel** and that **Mary**, **Bob**, and **Sam** are all **fascists**.

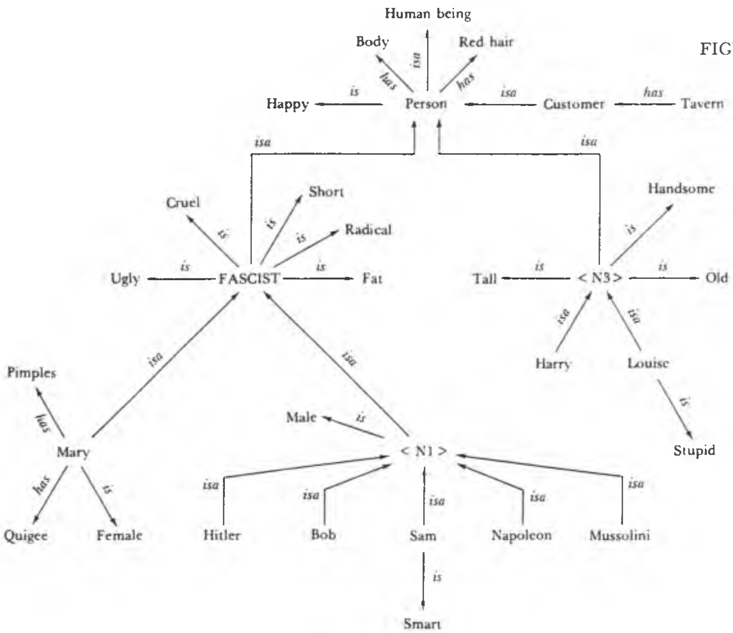

FIGURE 11-16

We have worked through some concrete examples to illustrate several procedures for rearranging the information stored in the data base. One procedure generalizes, deducing the common properties of concepts. Another makes special subdivision of classifications, realizing that specialized classifications can often be useful. Still another procedure looks for errors and inconsistencies and weeds them out of the system.

Exercise: From the data base for **pots** and **containers** (Figure 11-2), eliminate the inconsistency that states a **pot** is to be made of **metal**, but allows a **teapot** to be **ceramic**, even though it is defined to be a **pot**.

VIEWING THE DATA BASE

How much of the data can be seen at any one time by the interpretive process? So far, we have been drawing the networks so that everything is visible at once: It is not a difficult task to see just how things are interconnected. But it is quite possible that things are not so visible to the interpretive system. One way of thinking of this is to assume that the interpreter views the network by shining a flashlight at it. The only part visible is the part illuminated by the light. The question is, then, how wide is the beam of light produced by the flashlight?

Figures 11-17 & 11-18

We can see from the diagrams that there are many possible levels of visibility possible. (Note that the flashlight analogy is not completely accurate in that we show the network as getting more or less visible in terms of the number of arrows and concepts that can be seen, not in terms of physical diameter.)

The limitations in what the human retrieval process can "see" at any one time may be really a limitation of **short-term memory.** It is very likely that short-term memory holds the information on which the interpretive process is working. The capacity of short-term memory is measured in items, in psychological units. Now we can speculate about the nature of that unit. Perhaps a unit in short-term memory is a node. It is quite likely that the restricted number of nodes that can be retained in short-term memory may put some basic limitations on the ability of the interpretive processes to search out and evaluate information stored in the data base (long-term memory).

Try this thought experiment. Consider the incident at Luigi's, described

FIGURE 11-17

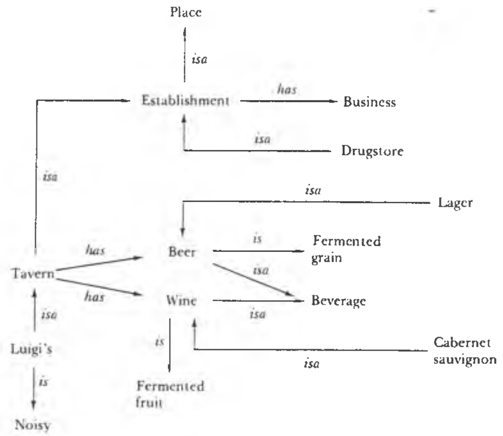

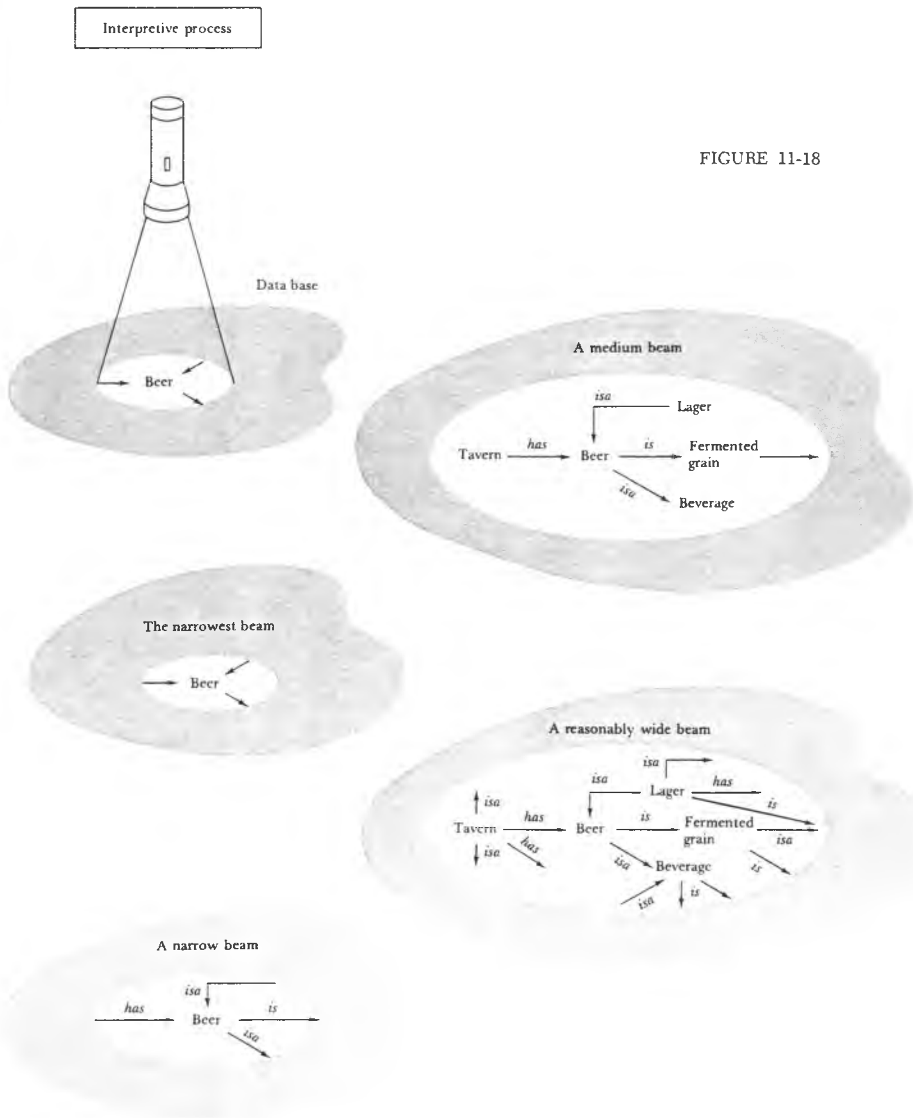
Interpretive process
Data base
Beer
A medium beam
isa
Lager
Tavern
has
Beer
is
Fermented grain
isa
Beverage
The narrowest beam
Beer
A reasonably wide beam
isa
isa
Lager
has
is
isa
Tavern
has
Beer
is
Fermented grain
isa
isa
has
isa
Beverage
is
isa
is
A narrow beam
isa
has
Beer
is
isa

FIGURE 11-18

in Figure 11-3 of this chapter and Figure 10-16 of the previous chapter. Imagine the scene. Luigi's is a dark, dim tavern, with customers sitting in booths in dimly lit corners. The owner, Al, is a friendly, personable chap. His dog, Henry, is always around. Suddenly, among a group of people in the corner, a scuffle erupts. Sam can be heard yelling at Mary. Henry, Al's dog, bites Sam, and somewhere in the melee Mary has hit Louise. Imagine the entire scene in your mind. Is the whole thing clear? If so, what kind of a dog is Henry? How long is his tail? Now look at his collar. What do the identification tags dangling from the collar say?

Most people find that as they imagine the scene, there is a limit to how much detail can be brought in at once. Originally, when they imagined the scuffle at the tavern, they claimed it all to be sharp and clear in their minds. Yet, when they are queried about the details of the dress of any of the participants, the color or length of their hair, or even the details of the tavern itself, they discover that it really is not quite so clear. When you are asked to examine the dog, the image of the dog fills your conscious awareness, and the rest of the incident, while still there in some sense, fades from its central location in the thought process. This can go on indefinitely. When you are asked to examine the dog in detail, he too turns out to be not so clearly noted. Examining his collar causes the rest of his body to fade from view. In fact, examining the tags hanging from his collar causes the collar itself to disappear into a haze.

It is tempting to argue that we can bring only a limited number of nodes into the short-term memory at one time (alternatively, into the illumination afforded by the flashlight). Thus, there is probably a central secondary node that refers to the entire incident at Luigi's (not shown in the diagram). This can be examined, but it is a general concept of the event, and it does not contain any details. When any of the details are followed, such as the node that represents any of the individual participants, then the other nodes that represent the details of the event are no longer visible. Although not shown in the diagram, presumably Henry points to a complex set of interrelations that define the appearance and exact details of his existence as a dog. When any of those nodes are examined with care, then the ones around it are no longer quite so visible, and the ones distant (such as Henry's role in the incident at Luigi's) are far removed from consideration.

UNDERSTANDING VERSUS REMEMBERING

One major implication of these ideas is that the memory system encodes the meaning of the material that has been experienced, not the material itself. Man attempts to understand rather than to remember.

Consider the following experiment (Bransford & Franks, 1971). Subjects are told a story made up of individual sentences.

The rock rolled down the mountain.

The rock crushed the hut.

The hut is at the river.

The hut is tiny.

When these sentences are actually presented, they are randomly interspersed with other sentences that are not part of the story, such as

The breeze is warm.

The ants ate the jelly.

The story is in the newspaper.

The jelly is sweet.

Thus, the four sentences that comprise one story must be extracted and put together from the entire set of sentences that actually are presented (and the other sentences can be combined to form four different stories). After each individual sentence is presented, the subject answers a simple question about it (Did what? Where?) to make sure that he indeed understands each sentence. Then, about 5 min after all the sentences have been presented, the subjects are presented with some test questions. They are asked to state whether they had actually heard these sentences before or not. For example, three possible test sentences are:

A. **The rock crushed the hut.**
B. **The rock crushed the tiny hut at the river.**
C. **The rock is tiny.**

Sentence A is, in fact, one of the original ones presented: The subject should state that he remembers it. Sentence C is not one of the originals; moreover, the meaning is different from that of any of the sentences. But sentence B is the most interesting one, for here the meaning is correct, but this particular sentence never was presented.

How do subjects remember these sentences? Here are the four sentences and their structure (Figure 11-19). Notice that for the way the memory system encodes information the four structures combine into one (Figure 11-19E). If this is all that is stored, how could you tell whether sentence B had been presented or not? Certainly, it is quite consistent with the diagram. Sentence C is clearly wrong: It contradicts what is stored. But sentences A and B are equally good.

FIGURE 11-19

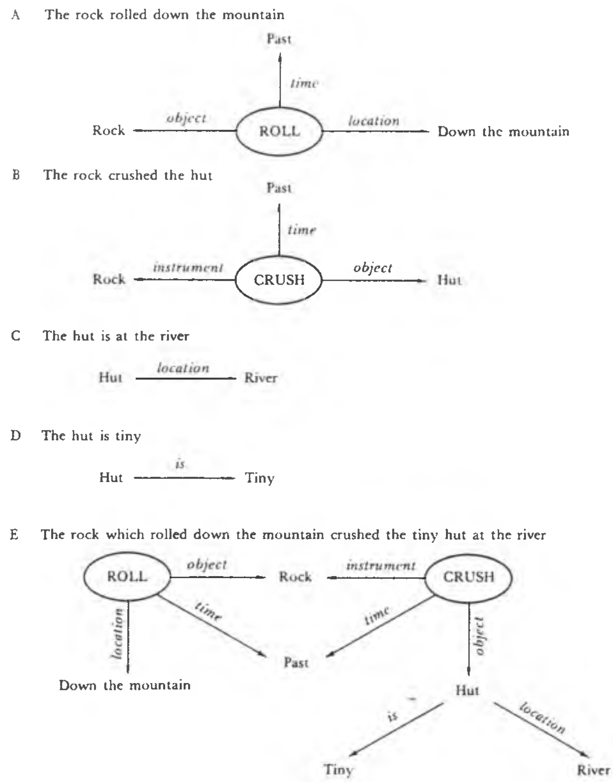

In the actual experiment, a set of test sentences was constructed with varying complexity: Some of the sentences contained a larger number of the original four ideas than others. Thus, consider these four examples:

Test 1: **The hut is at the river.**
Test 2: **The rock crushed the tiny hut.**
Test 3: **The rock crushed the tiny hut at the river.**
Test 4: **The rock which rolled down the mountain crushed the tiny hut at the river.**

Only the first of these test sentences, *Test 1*, actually was ever presented to the subjects. Yet, when subjects were asked to say whether they recognized any of the four sentences as one that was originally shown, they overwhelmingly selected *Test 4*, even though they had never in fact seen it. *Test 3* was next in terms of the number of times the subjects recognized it as occurring before. For both *Test 2* and *Test 1* (the only sentence actually presented) subjects usually denied having seen either of these sentences before. This is as it should be. The memory model encodes and remembers ideas and meanings. Hence, subjects should recognize sentences according to how well they fit the idea. Sentence *Test 4* is the most complete representation of the event as they remember it; sentence Test 1 overlaps only partially.

End of required part

MEMORY AS INFORMATION PROCESSING

The structure of memory described in this chapter presents a very different view of the nature of learning from that which results from studying how associations are built up in classical learning tasks. The difference is in the emphasis on the dynamic and integrative processes of memory. The system learns through an active interaction with its environment, rather than a passive build up of stimulus–response connections. This mode of operation pervades all of our experiences and is fundamental to the way we deal with the world we encounter.

Most of you probably learned the story of Hiawatha as a child. Can you remember what you thought when you first heard about him? Do you think your concept of Hiawatha has changed in the intervening years? If so, is it because you have reread the story or is it a result of learning more about the world, about Indians, about the nature of children's stories, and about the older American traditions of Indian folklore?

Classical learning theories typically rely on repeated exposure to specific stimuli as the basis of acquiring information. They have difficulties dealing with the fact that the understanding of a concept continues to be elaborated and embellished, even though the concept may never directly be encountered again. Such an evolution is a natural property of the type of memory system we have been examining. As more information about the world is accumulated, the memory system's understanding continues to grow and become elaborated. As an automatic by-product of this changing structure, our knowledge continually changes. Thus, it is very likely that your present recollection of the story of Hiawatha is quite different from the original. Any discussion about Hiawatha that might now be produced is determined in part by what was originally learned, but also by what has been thought about since that time. Try to recall it: Is not your recall primarily a reconstruction of what you believe the story must have been, rather than what it really was?

This continual evolution of the stored knowledge within the memory system has very profound effects on the way that new information is acquired. It suggests that there must be a tremendous difference between the way a message is encoded into a child's memory and the way the same information is encoded by an adult. For children, each concept encountered has to be built up from scratch. A great deal of rote learning must take place during the initial construction of the data base: Understanding is only slowly elaborated as properties are accumulated, as examples are learned, and as the class relations evolve. At first, most of the concepts in memory will only be partially defined and will not be well integrated with the other stored information.

Later in life, when a great deal of information has been accumulated and organized into a richly interconnected data base, learning should take on a different character. New things can be learned primarily by analogy to what is already known. The main problem becomes one of fitting a new concept into the preexisting memory structure: Once the right relationship has been established, the whole of past experience is automatically brought to bear on the interpretation and understanding of the new events.

For models of this type, the development of individual differences and idiosyncratic systems should be the rule, rather than the exception. Understanding evolves through a combination of the external evidence and the internal operations that manipulate and reorganize the incoming information. Two different memories would follow exactly the same path of development only if they received the identical inputs in the identical order and used identical procedures for organizing them. Thus, it is extremely unlikely that any two people will evolve exactly the same conceptual structure to represent the world they experience.

Be careful to note what is at the basis of this idiosyncratic development. We expect that both the basic structure of memory and the processes for manipulating and reorganizing information are similar from individual to individual. However, even though this basic machinery is the same, its operation will not necessarily generate the same memory products. What a person believes depends on what he has experienced and what sequence of inferences and deductions has been applied to the stored information. Even very subtle differences in the environment can produce different memory products, despite the fact that the underlying machinery for interpreting and remembering information may be common to all people.

The possibility that a basic set of processes can be used to deal with a variety of environmental contexts is, of course, a very adaptive feature of the memory system. But we might expect that the flexibility with which it can deal with new information would continuously change as the structure is built. It is seldom that an adult encounters an entirely novel event—one that is totally unrelated to his existing conceptual structure. Almost everything he experiences can be related to what he has encountered in the past. Even when he experiences clearly discrepant information, his conceptual structure is made up of such a complex and interdependent set of relationships that it resists revision. Thus, an adult is more likely to reject a discrepant input or change its meaning than to modify or change his beliefs. With children, the conceptual structure is not nearly so elaborate or so highly interconnected as that of adults. New experiences can be taken in stride, since contradictions seldom arise.

Perhaps the most interesting of the areas left unexplored is the interactive aspect of the human mind. People ask questions: They explore their own knowledge, they read, think, daydream and act. Even the most casual observations suggest that much of a child's behavior involves the engagement of his environment as he systematically seeks the information needed to build up his internal representation of the world. The model we have described here hints at these processes, but it does not do full justice to them. We have suggested some ways by which the memory system might ask for confirming evidence about the deductions and inferences it makes, but we have only scratched the surface of this very important area. The main problem at the moment is that there are no systematic tools for analyzing the natural exploratory behavior of people at work and play. A start has been made, however.

SUGGESTED READINGS

The concepts developed in Chapters 10 and 11 and representing the data base and its interpretation are novel. Little can be found today in the literature, but we predict that this will be changing rapidly for the study of long-term memory processes is increasing rapidly.

A symposium on the organization of memory held in Pittsburgh discussed many of these issues, so the interested reader might begin his search with that volume (Tulving & Donaldson, *Organization and Memory*, 1972). Perhaps the best place to start in that book is with our own chapter (Rumelhart, Lindsay, & Norman), for it presents a slightly more advanced version of the chapters in this book. Then, you could go on to the paper by Collins and Quillian, the chapter by Kintsch, and then the one by Bower. These chapters will provide the interested reader with a fairly comprehensive review of what has been done to this time, as well as specific literature references for further study. In addition a good review is provided by Frijda's survey article (to be published).

A number of experiments are now being performed on topics closely related to these issues. The classic work is the study on the reconstructive nature of remembering by Bartlett (1932). The studies referred to in this chapter were done by Bransford and Franks (1971), Collins and Quillian (1969), and Meyer (1970).

The model described in these chapters borrows heavily from the work on semantics by a number of modern linguists, but most especially the case grammar of Charles Fillmore (1968, 1969). The books in which the Fillmore articles appeared contain other papers highly relevant to the linguistics used here. The work by Winograd (1972) describes the development of a computer system for understanding language. This paper appeared in the journal, *Cognitive Psychology*, and current plans are for several other very important papers on the topic of language, memory, and thought to appear in this journal in the near future. They can't be referenced, because some of the papers we have in mind might not make it. Thus, you ought to skim through recent issues of the journal to see what new developments have occurred.

A number of relevant studies of memory are now beginning to appear in the computer science literature. Hence, the interested reader might wish to look at the book of collected papers by Minsky, *Semantic Information Processing* (1968). The book by Kolers and Eden (1968) also has some relevance, especially the chapter by Weizenbaum. The book by Loehlin (1968: paperback) offers a good introduction to models of personality somewhat relevant to the models of memory presented here. The Ph.D. thesis by Winograd (1970) (from which Figure 11-2 was taken) published in *Cognitive Psychology* (1972) is especially recommended, although it emphasizes language, not memory. In addition, Hunt's (1971) article, "What Kind of Computer is Man?" is a good article to read, although perhaps more important to the concept of man as an information processor than to specific notions of memory structure.