CHAPTER THREE

Comprehension in Skilled Reading

This chapter considers the main features of comprehension in skilled reading. At the outset, an important point must be emphasized. The distinction between lexical access and comprehension is not one between two stages of processing. The skilled reader does not first access a word and then comprehend. Instead, in normal reading comprehension and lexical access are concurrent ongoing processes. The reader who accesses a word generates both an updated model of the text and an immediate bit of local processing. This assumption indeed is necessary for the proposal of the last chapter that a text model provides semantic activation.

There are two major components to comprehension: *local processing* and *text modeling*. Both of these components depend on the reader's knowledge of word meanings, knowledge of domains related to the text content, and processes of inferring. To discuss these aspects of comprehension, we will consider the following simple story excerpt. Each sentence is numbered for reference.

(1) *Joe and his infant daughter were waiting for the doctor to get back from lunch.*
(2) *The room was warm and stuffy, so they opened the window.*
(3) *It didn't help much.*
(4) *Why wasn't there an air conditioner?*
(5) *Strange for this part of Manhattan.*
(6) *Suddenly the door swung open.*

In this brief excerpt there are six sentences and 47 words. It is a very simple text that could be read quickly and easily by a skilled reader. What are the processes of comprehension as this text is read?

LOCAL PROCESSING

By local processing, we mean those processes that construct elementary meaning units from the text over a relatively brief period. These include the semantic encoding of words and the assembling of propositions. These processes, at least the second one, take place within a limited-capacity processing system, which can process only a limited amount of text at any one time.

Semantic Encoding

Encoding the meaning of a single word is of course a lexical process and might as well have been discussed in a chapter on lexical processes as in one on comprehension. However, the meaning of a word is encoded in a way that is appropriate for its context. This is comprehension.

Actually, what a reader must have at the outset is an entry for each word in semantic memory. Such a semantic memory may be conceived as a network that links word concepts to abstracted meaning features and to other word concepts. The exact form of this representation is not important, although there certainly are a number of different systems (Anderson, 1976; Lindsay & Norman, 1977; Smith, Shoben, & Rips, 1974). In fact, the two main representation systems, semantic features and semantic networks, are formally nearly equivalent although they appear different on the surface.

Figure 3-1 shows a piece of a semantic memory system that would be activated during the reading of sentence (2). Its representation is something of a mix of current systems, being most similar to that of Lindsay and Norman (1977) except that it represents meanings directly (to the right of each concept) in addition to links to other word concepts.[1]

The reader has to have some word knowledge of this sort, however one chooses to represent it. The representation must provide sufficient structure so that while *hot* and *cold* are linked, the reader can say something like "*hot* is the opposite of *cold*" and not "*hot* means *cold*." He must also have knowledge that includes what is usually called *presupposition*. For example, to understand the meaning of *opened* in sentence (2), the reader presupposes that the object to be opened was closed. Such knowledge can be represented in different ways, perhaps best as an inference system that operates on the semantic network. (These matters have not received much attention in semantic network theories.)

Context-appropriate meanings

In sentence (2), the word *room* has more than one meaning, only one of which is represented in the network of Figure 3-1. It can also mean,

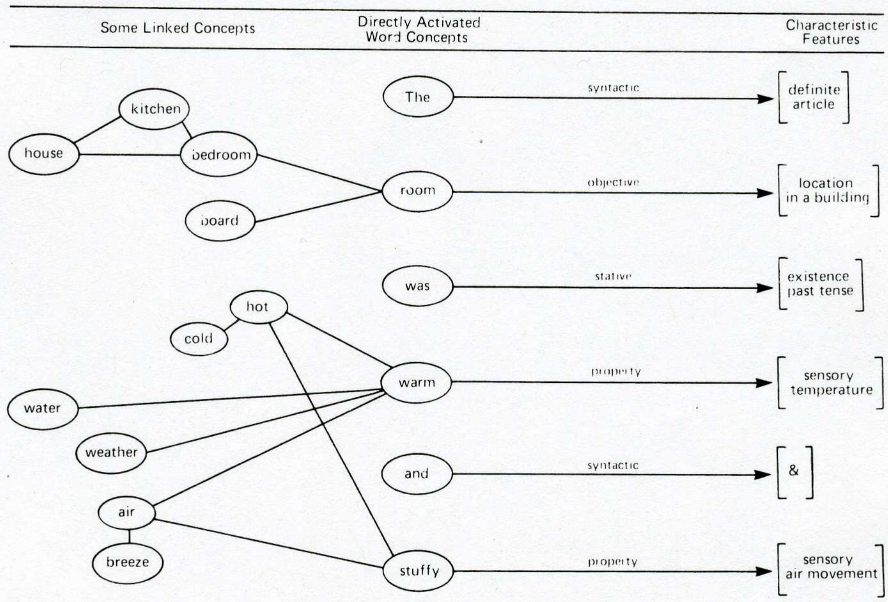

Figure 3 1

Figure 3-1. A piece of semantic memory network activated by *The room was warm and stuffy*. Word concepts are nodes in the memory network. They are closely linked with meaning features that "define" the concept. (This is a departure from network representations such as Lindsay & Norman's [1977], in which labeled links exist only among concept nodes.) Meaning features are connected by links that label the general class of each concept. These include *syntactic modulators*, *actions*, *states*, *properties*, *objectives*, *abstractions*, and *names*. To the left are links to other concepts that potentially receive activation.

roughly, *open space without restriction*, as in "The university has no more room for expansion." It also has common metaphorical extension, as in "room to grow." Such multiple meanings are very common in natural languages. Furthermore, the multiple meanings of a word are not always so closely related as in the room example. The word *bar*, for instance, has two meanings with little in common, as both lawyers and bartenders will attest.

How does the reader encode the word in its appropriate meaning? The intuitive answer is that context determines this process. When one is reading about an attorney's struggles with an ethics committee, *bar* is encoded as [lawyers' guild]. When one is reading about the best places to have a pina colada in Philadelphia, *bar* is encoded as [drinking place].

This is undoubtedly true, but when does context do its work? It would be convenient if there were always powerful semantic cues before the word is encountered. But what about when the word is inconsiderate enough to occur early in a sentence? *Room* in sentence (2) is such an example. Often, there is also context assistance from the previous sentence, but not always: *It was a beautiful day. The fair. . .* At just this point how will *fair* be encoded? Certainly the remaining context will force an appropriate encoding of *fair*. But at the time of lexical access, there are many possibilities that fit the context. *The fair weather has been with us for a week. The fair Gwendolyn would be mine by nightfall.*

The solution is in the activation model. As context accrues, semantic activation rises and falls. If context heavily biases one meaning over another, then that appropriate meaning will have an activation advantage. However, because activation spreads, even inappropriate meanings will receive some activation. So two things should happen: (1) All meanings should be activated, and (2) the appropriate meaning "wins out" strictly as a function of the rate of its activation and the rate of decay of activation for the inappropriate meaning.

The first of these predictions, counterintuitive as it is, seems to have been confirmed. Swinney (1979) had subjects perform a lexical decision task, i.e., to decide whether a letter string was a word or a nonword. The lexical decision was made visually just at the moment that an ambiguous word was heard through a headset. It is known that lexical decisions are affected by prior semantic activation. Thus *butter* preceded by *bread* is accessed more quickly than *butter* alone, and the lexical decision is faster, relative to a control (Meyer & Schvaneveldt, 1971). The question is whether such semantic activation would include both meanings of an ambiguous word or just its contextually appropriate meaning. For example, subjects had to make a lexical decision for either *ant* or *spy* upon hearing the word bug. There are two meanings of bug, one related to insects and one to surreptitious listening devices. In Swinney's experiment the subject heard the sentence "The man was not surprised when he found several spiders, roaches, and other *bugs*." Just as *bugs* was heard, the subject saw either *ant* or *spy* or a control word, *sew*. If context selects the appropriate meaning of *bug* and if activation spreads only to concepts related to the appropriate meaning, then *ant* should be more quickly accessed than the control word, but *spy* should not. But in the experiment *both* words were accessed faster than the control. Even the word not appropriate for the contextually selected meaning was activated.

The principle that emerges is that all meanings are activated, at least for a brief time. There is also some support for the second implication of the activation theory, namely that the contextually appropriate meaning will "win out" as its activation continues while the inappropriate one decays. When Swinney (1979) delayed the lexical decision by four syllables—i.e., four spoken syllables elapsed between the spoken ambiguous word and the printed lexical decision word—only the word related to the appropriate meaning (*ant* in the example) was processed more quickly. Also, using a different task, Tanenhaus, Lieman, and Seidenberg (1979) found evidence that both meanings of a word were available immediately but after 200 milliseconds only the appropriate meaning was available. It remains possible that if context very heavily biases one meaning over another, even early activation becomes more selective (Simpson, 1981). Overall, however, the conclusion is that semantic encoding proceeds in two phases. In the first phase, lexical access activates all meanings associated with a word. The appropriate meaning receives more activation and is selected for encoding.

This second phase, encoding the appropriate meaning after activation, may be influenced by an ordering principle built into the semantic representation. The model of Hogaboam and Perfetti (1975) assumes that more dominant meanings are selected before less dominant meanings to be tested against the encoded meaning of the context. Their evidence indicates that the dominant (most frequent) meaning of a word is more readily available to the reader than is the subordinate (less frequent) meaning. This ordering assumption can be incorporated into an activation model. We assume that the resting level of activation is higher for the more frequent meaning. Lexical access does affect the activation level of all meanings, but when the alternative meaning is appropriate, the dominant meaning will remain activated for longer than the subordinate meaning.

Propositional Encoding

The encoding of individual words enables the encoding of elementary text units, or propositions. We return to sentence (2) of the story excerpt to illustrate. The propositions of that sentence, *The room was warm and stuffy, so they opened the window*, can be represented as follows:

1. exists (room)
2. warm (room)
3. stuffy (room)
4. exists (window)
5. open (they, window)
6. because (5[2 & 3])

Thus there are six elementary propositions in this sentence. Two, however, are existential propositions that merely assert the existence of some referent. These propositions, in a sense, answer the questions "what room?" and "what window?" that would arise in propositions 2 and 3. No existence proposition is required for *they* because it is anaphoric. It refers to something whose existence has already been established (in the first sentence of the story).

The central element in a proposition is the predicate. Grammatically, it may be a verb, an adjective, or certain conjunctions. It predicates a relationship between two or more concepts in a sentence, as in proposition 5, or it predicates some property or state of a concept, as in propositions 2 and 3. (The existential predicate "to be" can be said to predicate the property of existence.) Proposition 6 is an *embedding* proposition, in that it combines two propositions by embedding them into a third one. It predicates a causal relationship between proposition 5 and propositions 2 and 3.

In short, *propositions* are abstract, elementary meaning units that comprise the meaning of a sentence. A text may be defined as a set of coherently related sentences and may therefore be represented as a list of interrelated propositions. The major theory of text representation based on this assumption is described in Kintsch (1974) and Kintsch and van Dijk (1978).

Working memory

On encountering the words in a sentence and encoding them, the reader assembles them into propositions. Or rather, propositions represent the meaning information that the reader assembles from a sentence. This assembly occurs within a limited-capacity processing mechanism, the same mechanism that can be applied to remembering a string of digits or multiplying pairs of two-digit numbers. Since this working memory (Baddeley & Hitch, 1974; Newell & Simon, 1972) is a limited-capacity mechanism, it follows that the amount of information it can handle is limited. We can assume, with Kintsch and van Dijk (1978), that the reader can hold only a few propositions in working memory at one time. As new propositions are assembled, previously assembled ones are vulnerable to memory loss. The trick, in general, is to quickly integrate the assembled propositions into a representation that can survive in long-term memory. If a sentence is very long, and especially if its end part has a proposition that needs to be integrated with its front part, extra processing effort must be applied.[2]

Thus, reading as a local process is a matter of assembling and integrating propositions for longer term memory. One fact that gives some support to this account is that the time it takes to silently read a short text is a function of the number of propositions, even when number of words is controlled (Kintsch & Keenan, 1973). There are many other important aspects of a model that describes reading in terms of propositions, including exactly how memory is limited, how a representation of the text is built up in memory, etc. These are discussed in Kintsch and van Dijk (1978). Although many processing details are not well understood, a central point is that the reader processes some kind of text meaning unit that seems reasonably well approximated by the proposition. In the course of this processing, he constructs, in the ideal case, a coherent set of propositions that constitute the text's meaning.

Integrative processes

Processes that integrate text material are a continuous part of reading. As a local process, *integration* refers to combining successively occurring propositions with each other. For example, in the visit-to-the-doctor story, the six propositions of sentence (2), are integrated with each other and with the representation of sentence (1). The latter we assume is at least temporarily in short-term memory. Thus *local* integrative processes are those that occur by linking recently formed representations with new propositions as they are encoded.

A major mechanism for local integration is the recurrence of an element in different propositions. Kintsch and van Dijk (1978) refer to this recurrence as *argument overlap* (nouns are "arguments"). For example, the two propositions (2 and 3) in sentence (2) both contain *room*. Accordingly they are quickly integrated in working memory: roughly speaking, the room that is warm is also the room that is stuffy. A somewhat more interesting example can be seen in proposition 5 of sentence (2), *they opened the window*. The integration here has to occur by linking the pronoun *they* to something earlier. The reader has to link *they* to *Joe* and *his infant daughter*, which are represented in a proposition of sentence (1).

Such a linkage would be carried out first by recognizing that *they* is anaphoric, i.e., having some antecedent element as its reference. Then a search of memory is made for the reference, and *Joe* and *his infant daughter* are quickly found. All this happens quite quickly and automatically unless the antecedent is no longer in short-term memory. Thus local integration is a process that depends on linguistic signals that trigger attachments in memory.

In addition to pronouns, another type of linguistic trigger is the definite article. It tells the reader either to expect a match in memory or to build one. For example, in sentence (2) the reader immediately encounters "the window." But what window? His text memory contains no window, but the definite article *the* is a signal to infer the existence of a window. So proposition 4 establishes the existence of a window. In order to do such inferring, the reader must have knowledge about rooms, including the knowledge that they often have windows. The previous mention of *room* has activated the knowledge needed so that a new proposition may refer to windows, floors, furniture, and any of the things that rooms ordinarily have. (This is schema activation, discussed in the next section.)

Of course, a very useful linguistic trigger is the repetition of a word itself. The next mention of "Joe" or "the doctor" or "the window" in the story will immediately lead to a link with a previous proposition containing one of those words. However, the word itself is not the key. Haviland and Clark (1974) demonstrated this with sentences of the following kind: (1) *Marvin liked beer.* (2) *The beer was warm.* The subject who reads these sentences encounters a repetition of the word *beer*. However, the first appearance of the word has not established the existence of a particular reference. (There would be no existential proposition representing the existence of beer.) Thus the linking attempt for the second occurrence of *beer* would fail. Haviland and Clark (1974) found that subjects took longer to read sentence (2) than when it followed a sentence that established reference by mentioning some specific beer.

An important process that occurs as part of local integration is *reinstatement* (Lesgold, Roth, & Curtis, 1979). It is fairly typical for some proposition to be displaced from working memory by subsequent text processing. As we noted, one thing that can happen is an active search of memory for the needed information. Sometimes the information needed is fairly quickly reactivated (or "reinstated"), because it is kept in the foreground by the text that follows. For example, Lesgold et al. compared the time to read a given sentence under several conditions. In one, a necessary antecedent was kept in the foreground. In the other, it slipped to the background. For example, this text mentions *forest* in the first sentence. Then it keeps the forest concept in the foreground, even though it is not referred to again:

> *A thick cloud of smoke hung over the forest.* The smoke was thick and black and began to fill the clear sky. Up ahead Carol could see a ranger directing traffic to slow down. *The forest was on fire.*

In the comparison condition the intervening sentences (those not italicized) did not refer to *smoke* nor to *forest ranger*, but to activities inside the car. (What car? The reader's inference is invited.) The difference between keeping the forest in the foreground and letting it slip into the background is that the final sentence (italicized) is read more quickly in the foregrounding condition.

Thus, as Lesgold et al. describe it, there are three integrative processes that can occur when a sentence is read: (1) An immediate match can be found in active memory (or short-term memory), based on recently encoded propositions. (2) When the activated memory does not contain the match, the information is easily *reinstated* when intervening text has kept it foregrounded. (3) When there is no immediate match in active memory and when the required information has not been found, a search of memory is required. Comprehension in such a case will depend on the reader finding the relevant information and perhaps making a "bridging inference" (Haviland & Clark, 1974).

Thus integrative processes depend on linguistic triggers and the accessibility of linking propositions in memory. To the extent that what is accessed is more than the most recently encoded text, we are discussing not merely local processing but the construction by the reader of a model of the text.

THE READER'S TEXT MODEL

There is, of course, more to reading a text than encoding words and propositions. The reader encodes these propositions in the context of knowledge about concepts, knowledge about inferences (inference rules), knowledge about the forms of texts, and general knowledge about the everyday world. By *text modeling* we mean the processes by which the reader combines such knowledge with local processes to form a representation of the text meaning. It is this representation that the reader consults at some later time to recall or to answer questions about what has been read.

Knowledge and Inference

Consider again the story excerpt.

(1) *Joe and his infant daughter were waiting for the doctor to get back from lunch.*
(2) *The room was warm and stuffy, so they opened the window.*
(3) *It didn't help much.*
(4) *Why wasn't there an air conditioner?*
(5) *Strange, for this part of Manhattan.*
(6) *Suddenly the door swung open.*

After reading this excerpt, a reader might be asked to recall it exactly. It is possible that his recall attempt would include the following in place of sentence (2): "Joe opened the window in the office because it was hot and stuffy." The chances are one has to check back with the original just to be sure of what the differences are. First, *hot* is recalled instead of *warm*. This is a local text change based on encoded word meaning. Such a change can be accounted for by reference to semantic memory connections and lack of precision in encoding. However, recalling *office* instead of *room* is a different matter, although it looks quite the same. *Office* and *room* are not so closely related in semantic memory. Instead, the error has occurred because of what the reader knows about the everyday world, i.e., that people visit doctors in places called "offices." (If the text had been about Joe's waiting in a room for a late-arriving airplane, the *office* error would not have occurred.) Finally, the same sort of process causes the reader to recall *Joe* instead of *they*. The reader knows that *infants* are not plausible agents for opening windows and infers that the real agent was the adult. Thus some text-recalls reflect processes of word encoding while others reflect a combination of what the reader knows and what is readily inferred.

Some of the inferences a reader makes are semantically or logically implied by the text propositions (entailment) while others are not. The office inference examples above are in the latter category. Their plausibility is strictly psychological, based on the reader's world knowledge. This means, of course, that they are not guaranteed to be correct inferences. In fact, Joe could well have been in a "waiting room" rather than an office. The *Joe* for *they* case is trickier. It is in fact a logical entailment that if Joe and his daughter *(they)* opened the window, then Joe opened the window. So it is a necessarily true inference. But such recalls, if the reader really remembered *they*, would violate pragmatic conventions, observed by the writer, on saying not only what we know to be true but what we know will not be misunderstood.[3]

For a better example of a logically or semantically forced inference, suppose that a reader recalls something about the window being closed when Joe and his daughter entered the room. It is semantically forced because sentence (2) tells of opening the window. Something can be said to have been opened only if its initial state was closed.

Despite their differences, it is probable that all these inferential alterations are governed by a single comprehension principle: In processing a text, the explicit propositions are combined with the reader's knowledge to produce a text model. This knowledge includes what the reader knows about word meanings, e.g., the relationship between *open* and *closed* and *hot* and *warm* as well as what he knows about doctor's offices and other matters of the world. The organized knowledge about such things is a *schema* (Anderson, Spiro, & Anderson, 1978; Bartlett, 1932; Rumelhart & Ortony, 1977).

Schemata

A *schema* is a conceptual abstraction containing slots (or variables) to be instantiated in various ways; it can apply to simple word concepts as well as doctor's offices and restaurants (a favored example of schema theorists). For example, a schema for window opening would include slots for the manner of opening—by lifting, by pushing, by turning a handle, etc. The core of the schema, its invariant part, has to do with causing a window which is closed to become open. More complex schemata have the same properties.

A visit to a doctor includes variable slots relating to appointments (essential or advisable?), receptionists (is the receptionist also a nurse?), and waiting rooms (*Reader's Digest* or *U.S. News and World Report?*).

In the previous discussion the examples of hypothetical "errors" in recall were partly to demonstrate how normal comprehension works. As we realize that the reader's knowledge about many things is important in comprehension, our theory of reading must take account of such knowledge. We must assume that there is available to the processing mechanism the full range of the reader's knowledge—knowledge of semantic relations, rules for inference, and everyday general knowledge. It is possible to suggest that much relevant knowledge is in the form of schemata, organized knowledge structures. Even some inference rules can be defined, at least partly, as the application of schema knowledge to texts which, as all texts must, leave out some of the information contained in a schema.

An experiment by Anderson, Spiro, and Anderson (1978) demonstrates how schemata organize information for a reader. Subjects read either a passage about having dinner at a fancy restaurant or one about a trip to a supermarket. Eighteen key food items appeared in each of the two passages. The question was whether these 18 items would be recalled equally well by readers of the two passages. The schema hypothesis predicts otherwise, on the assumption that a restaurant schema—or "script" as Schank and Abelson (1977) refer to it—organizes food items more powerfully. For example, the restaurant schema includes an ordering principle for classes of foods (appetizers, soups, salads, entrees, desserts). Restaurant meals also contain fewer food items than supermarkets. In a supermarket, oregano is organized at the same level as ice cream. In a restaurant meal it is not. For at least these reasons, the restaurant schema should enable more recall of food items than the supermarket schema, and this was in fact demonstrated by the experiment.

The contribution of a schema in such a situation is at least partly as a plan for retrieval. Whether a schema is used during comprehension as well as during recall requires a different demonstration. Two examples: Bransford and Johnson (1973) constructed extremely vague passages that were nearly unrecallable by subjects. Their vagueness meant that constructing a text model was very difficult. However, when subjects were given a picture clue or a title prior to reading the passage, their recall was much better. Subjects also rated the passages as more comprehensible.

A second example comes from Dooling and Lachman (1971), who constructed a metaphorical passage in which the hero asserted that "an egg, not a table, correctly typifies this unexplored planet." When subjects were given the title "Christopher Columbus" *before* reading the passage they recalled it quite well. When the title was given *afterward*, it was not much help. Thus, such studies demonstrate that schemata are more than plans for retrieval. They are also structures which serve the comprehension of a text. (See Figure 3-2 for an example of how schemata are important in comprehension. It is a genuine example, not one for experimental purposes.)

Exactly how schemata facilitate comprehension is not quite so clear. As Anderson et al. (1978) put it, a schema serves as a scaffolding on which to construct the meaning of the text. However, there are several things that must happen to make this work. First, the appropriate schema must be activated (or identified). Then the text interpretation is guided by knowledge activated by the schema. For example, the schema directs the reader to understand the passage of Dooling and Lachman as about Christopher Columbus and not about an egg. In addition, as we have seen from the story about the doctor's office, there are inferences to be made and the appropriate schema directs these as well. (These functions of schemata are discussed by Rumelhart & Ortony, 1977; Schank & Abelson, 1977; and others.)

However, there may be equally powerful schematic functions regarding the reader's expectations. Refer back to the doctor's office story. There are only six sentences, but already expectations have been aroused by the doctor's office schema. The reader expects perhaps an examination (of the daughter?) by the doctor. The reader at least expects the doctor to show up and do something with Joe and his daughter. If something in the story does not fit the schema, it may receive extra attention. In a sense, the activation of a powerful schema and a text that conforms closely to the schema at all points may be processed mainly in a top-down manner. The reader can simply confirm what he expects. (This may make for rather uninspiring reading in some cases.)

What happens to information that does not fit the schema? It is not simply filtered out. True, people often recall the gist of a passage, leaving out detail. But they have very good recognition memories for actions not typical for the schema (Bower, Black, & Turner, 1979; Graesser, Gordon, & Sawyer, 1979). Even recall, under certain conditions, shows good memory for atypical schema information (Graesser, 1981). Thus a model needs to account for the reader's better recall of gist, his poorer recall of detail, and his ability to remember atypical information.

LOS ANGELES — The latest fallout of the space program is an astonishing data-recording system developed by scientist E.T. Seti.

The brain of the machine, located on one end, is known as the Data Stream Imprint System, or DSIS, a NASA-copyrighted graphite linear feed designed to perform a message-recording function. Protection against the hostile environment of space is provided by a cellulose-fiber, reinforced-resin protection layer.

The opposite end of the machine incorporates an ingenious solution to error-correction problems, an abrasive data character erase module. In the erase mode, the module is briskly rubbed across the characters to be deleted and, by the phenomenon known in physics as "lift-off," the undesired characters may easily be expunged.

Reliability tests, conducted under rigid National Aeronautics and Space Administration test parameters, revealed an extremely low failure rate, with the graphite fracturing only once in every 1,000 performances. (A peripheral product has been devised, to be used in the rare event of such a failure, that incorporates the use of a clever device called the Linear Feed Maintenance System. The LFMS literally "sharpens" the recording implement by removing the cellulose-fiber protection layer, exposing a fresh length of graphite.)

When I called Dr. Seti to ask him about the instrument, he conceded that each time it is "sharpened" it shortens, which limits its effective half-life to the operator's ability to handle short stubs. That problem is elegantly solved, however, by the interchangeable "throwaway" concept, enabling the operator to select, at his option, the manual override mode provided. This allows him to discard the stub and replace it with a brand new data-recording instrument from the box of spares provided.

These devices cost NASA $237.50 each. They will be offered to the public at 49 cents, including an instruction booklet, a fully equipped LFMS and a 90-day warranty.

Another spin-off of the space program, soon to be made available to the public, is a portable unitized earthwork synthesis system, sometimes referred to as the terrestrial transport shuttle. This product is designed to relocate dirt among piles, or even dig a hole, when properly manipulated by the operator.

A remarkable design feature is the control stick, better known to NASA as the air-to-ground interface contour adjustment probe, developed by engineering physiologists for optimal utilization by low-skill-level personnel. The stick permits the operator to assume a command and control function over the leveraged tactile-feedback geomass delivery module located on the Earth end of the system.

An ergonomic task analysis, conducted under simulated field conditions and timed by stopwatch, was carried out by an operations research team. As a result, a simple step-by-step program was devised, which can be used by any nonscientist in any interactive man-machine situation requiring the relocation of dirt from one pile to another.

Because the proper use of the device requires a "shoving" action, NASA marketing experts first recommended that the gadget be called a shovel, but later rejected the term as insufficiently descriptive.

A recent inventory of NASA's warehouse uncovered an overrun of nearly 3 million terrestrial transport shuttles on hand, for which Uncle Sam shelled out 1,200 disbursement dollars apiece. The shuttles will be offered to the public on a first-come first-served basis, at a surplus-disposal price of $7.37 each, or two for $9.39 — no refunds, no exchanges, limit two to a customer.

There is still another NASA product of which the limitation of space permits only brief description — the individualized reciprocating bicuspid plaque-level limiter, which will be released to the public under the simple Air Force designation of the unitized chopper disequilibriator.

This sophisticated machine was designed to be used by astronauts during zero-gravity teeth-cleaning missions. A preliminary parametric analysis has indicated that it operates most efficiently when used in conjunction with a viscous detergeable purging agent that the astronaut squeezes out of a compressible tube. Tentative strategy calls for this operation to be followed by a flush-out procedure that involves the use of large drafts of an H_2O dilutant that will be specially created for a weightless and benign environment.

Pricing information on this product will be released at a later date. More data on any of these products may be obtained by contacting NASA's Washington office. Dealers are requested to submit sealed bids for quantity purchases.

Figure 3-2. This excerpt from a newspaper column shows that schemata can be necessary for understanding even "real life" texts, as opposed to those made up by researchers. This one was written by Jack Catran for the *Los Angeles Times* and was reprinted in the international edition of the *New York Herald Tribune*, from which it is excerpted here with permission.

One model that does this has been proposed by Graesser (1981). (Similar models have been proposed by Bower et al., 1979; Norman and Bobrow, 1979; Spiro, 1980; and Thorndyke and Hayes-Roth, 1979.) This model assumes that the memory representation includes a "pointer" to an appropriate schema and a set of tags for elements that are atypical for the schema. Tags are also attached to typical information, but the more typical it is, the less likely it is to be tagged. If this seems to imply that readers will fail to remember the typical schema information, that is true, in a sense. They recall such information very well because it is part of the schema that is activated. They sometimes recall it when it is not actually there. Thus memories for texts (as well as other things) will be a blend of more highly encoded information and more highly inferred information. Both types of information may be part of the reader's text model.

An Example of a Reader's Text Model

Figures 3-3 and 3-4 illustrate the kind of text model that a reader might construct for the doctor's visit story. Figure 3-3 shows successive lines of the text going down and processing activities going from left to right. The processing activities are partly sequential but overlapping. Thus, as the text words are read, propositions are encoded, schemata are activated, and a text model is constructed and updated. Both propositions and schemata are used to construct the text model.

Figure 3-4 shows how the text model might be updated over successive sections of the story. The representation is partly hierarchical, in that the most central component *at the moment* is represented at the top level. While other parts of the model might be hierarchical, they are represented as simple network connections. However, the network is structured also by labeled links. The links correspond to cognitive structures required to understand the content of a discourse. In the case of a story, the underlying cognitive structures are sequences of events that are based on goal-directed activities. Thus *characters*, their *plans*, *obstacles* to their plans, their *actions*, and *consequences* of their actions constitute the central event structure of a story. There are a number of proposals for how to describe these structures (Graesser, 1981; Lichtenstein & Brewer, 1980; Omanson, 1982a; Glowalla & Colonius, 1983; Warren, Nicholas, & Trabasso, 1979). They differ in a number of ways, but share the basic assumption that the structure of events is the central underlying description for narratives.[4] This is what Figure 3-4 represents, although its representation of two-level structure is not typical.

The successive text models of Figure 3-4 can be summarized as follows. After encoding the propositions of the first sentence and having schemata appropriately activated, the reader's text model is that there is

Text Sentence Read →	Propositions Encoded →	Schemata Activated →	Text Model Constructed
1. Joe and his infant daughter were waiting for the doctor to get back from lunch.	1. *Joe* 2. *daughter* 3. *infant daughter* 4. *doctor* 5. *wait (6)* 6. *return doctor* 7. *lunch doctor*	a. *father–daughter schema* b. *doctor-visit schema* c. *doctor-lunch schema*	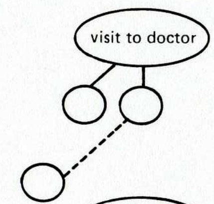
2. The room was warm and stuffy, so they opened the window.	8. *room* 9. *warm room* 10. *stuffy room* 11. *window* 12. *open window* 13. *cause (12, 8, & 9)*	d. *window–opening schema*	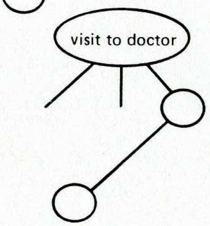
3. It didn't help much.	14. *not help (12)*		
4. Why wasn't there an air conditioner?	15. *no air conditioner* 16. *why 15?*	e. *consciousness–of–character convention*	(see figure 3 - 4)
5. Strange for this part of Manhattan.	17. *strange (15 & 18)* 18. *Manhattan*	f. *Manhattan schema*	
6. Suddenly the door swung open.	19. *open door* 20. *suddenly (19)*	g. *room-entering schema*	

Figure 3-3. Some parts of the processing of six sentences from the doctor's visit story. Propositions are encoded, schemata activated, and tentative text models constructed. (Propositions are represented in a simplified manner.) This is the order in which these processes occur.

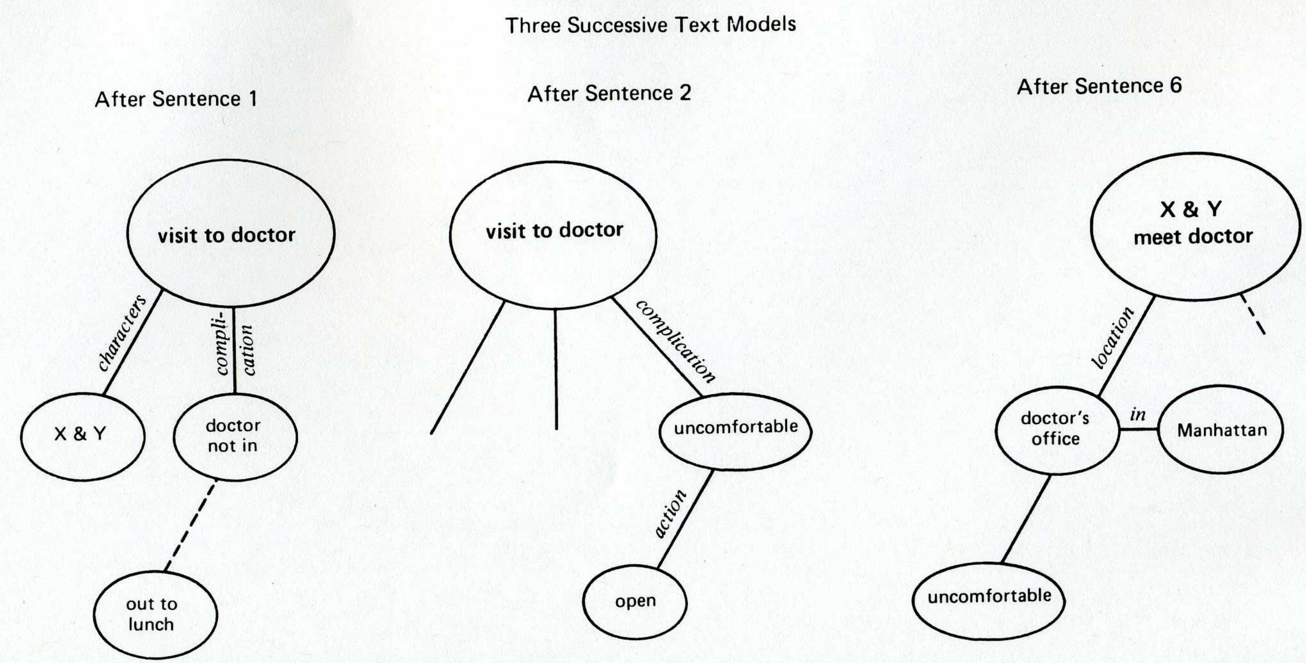

Figure 3-4. Three successive text models in schematic representation. After sentence (2), there is new information, and its potential importance adds it to the Visit head node. The previous links to this node are suppressed for the moment but still available. After sentences (3) and (4), there would be minor changes to reflect the consequence of the action. After sentence (5) (not shown), a location information would be added. After sentence (6) a new model is needed, or soon will be, to accommodate the door opening.

a VISIT TO THE DOCTOR by two people whose goal is blocked for the moment by the doctor's absence. The reader's model is of course very tentative. Encoding the second sentence produces the lower level change. The top-level event is still VISIT TO THE DOCTOR, but the model now represents a second minor complication, the discomfort of the room. Following the propositions of the third sentence, the text model changes only slightly to reflect no consequence of the action taken. Similarly changes following sentences (4) and (5) are minor. They keep the VISIT TO THE DOCTOR as central and add lower level structure.[5] (These models are not shown in Figure 3-4.) The final text model shown is one that the reader might construct following the last sentence. At this point the reader is ready for a switch in event dominance, and VISIT TO THE DOCTOR and MEETING WITH X are now both at the top of the structure. The room-entering is a new event with potential importance. It will quickly recede in importance with more text, either because it was of no consequence (e.g., a mistaken entry) or because its consequence begins to dominate the model (if, say, it is the doctor and he announces a grave diagnosis).

As the story progresses the text model changes, and the *kind* of structures that dominate the model may also change. For example, the *plans* of the characters should begin to dominate their actions and provide the highest level of the text model (Lichtenstein & Brewer, 1980).

Discourse Types

So far, we have focused on knowledge about word meanings and everyday knowledge to explain how a reader constructs a text model. Another relevant sort of knowledge is the kind of text being read. The doctor's visit story is one kind of text type, a narrative. The reader's text model shown in Figures 3-3 and 3-4, that is, one based on event structures, is appropriate for stories. (It is also appropriate for ordinary perception, but that is another matter). In fact, such an event structure is a schema for stories. There are descriptions of story schemata that are complementary to the one described here. These are the story grammars (Rumelhart, 1975; Mandler & Johnson, 1977; Stein & Glenn, 1979; Thorndyke, 1977). A story grammar is a rule-based description of regularities in the structure of stories, predicated on the fact that stories conventionally have settings and episodes. In fact, their structural descriptions are not so different from the kinds of categories appropriate to describe the structure of events. After all, a story is a text about some events that have certain properties.

Exactly what these properties are is an interesting question. Stories do seem to require some aspect of conflict and striving, not a mere haphazard chain of events. However, these aspects can be well captured by event structures that include at least hierarchical properties that subordinate the events to something, especially perhaps a character's plan for action, as Lichtenstein and Brewer (1980) argue. To the extent that story grammars capture features unique to stories as text forms, as opposed to stories as event sequences, there would be an additional source of schematic knowledge for the reader.[6]

Of course, stories are not the only discourse type. There are scientific texts, descriptive newspaper accounts, political essays, letters, etc. The structures of these texts may provide the reader with schemata concerning the distribution of information (Meyer, 1975). For example, newspaper accounts have strong conventions about providing specific information regarding the place, time, and persons involved in a traffic accident or a hold-up. A reader can potentially use such information at least to guide an information search. On the other hand, there is little evidence confirming the use of specific-discourse-type schemata. One exception is scientific journal texts, where there is some evidence that organizational features of those texts are used by readers (Vesonder, 1979).

In addition, there is evidence that readers are sensitive to the linguistic style of a text, which is in fact a matter of discourse type. A legal style, an academic style, and a biblical style are examples that readers can recognize readily. Brewer and Hay (1982) report that subjects' recall of a text reflects the style it was written in, independently of content. For texts written in mixed styles (some parts standard and other parts biblical), recall of standard parts was often in biblical style and vice-versa. Thus it appears that the reader's model of a text will include information about the type of discourse as well as about the content.

SUMMARY

Comprehension in skilled reading includes a number of local processes and text-modeling processes. Local processes are those that operate as the reader gets some meaning out of sentences. They include the encoding of word meanings and the assembly and integration of propositions. Some of these processes will be limited by the functional capacity of working memory. Text-modeling processes combine higher level knowledge and inference processes with the output of local processes to construct a text model.

The encoding of word meanings depends on the highly structured representation of word meanings in memory. The encoding of appropriate meanings is determined by context, but there is a general activation process that briefly affects even meanings not needed by context. As word meanings are encoded, they are assembled into propositions in working memory. There are limits to the amount of information held by this system. Integration of propositions occurs both within sentences and across sentences. Integration is triggered by linguistic devices and can occur as a result of immediate memory matches or reinstatements of deactivated memories or by bridging inferences.

The reader constructs a text model by applying schematic knowledge and inference processes to the results of local processing. These schemata include a wide range of conceptual structures that depend on everyday knowledge, word meanings, and discourse types. The reader's text model is continuously updated by the local processes and the reader's knowledge and expectations.

NOTES

1. Another difference is the basic semantic categories used. These are the labels (above the arrows) to the meaning features in Figure 3-1. There are seven such categories that do not correspond exactly to categories of case grammar nor to any other system. For example, nouns that refer to physically definable objects are *objective*. Their status as locations, human agents, etc., constitutes part of their meaning features. Nouns may also denote *abstraction* (e.g., "understanding," "justice"), *property*, *action*, and *state*. Syntactic modulators are words that are essentially logical operators on concepts, especially the so-called function words.
2. This characterization is fairly general to avoid problems of technical detail. In fact, Kintsch and van Dijk (1978) assume that propositions are immediately chunked into units the size of which is governed by memory limitations and text features. Each unit is a processing cycle. Thus there is already some assembling and integrating, as I have called them. The number of these units that are held simultaneously in memory is a model parameter, usually set at two. As new propositions are encoded they are attached to propositions stored in memory on the basis of argument overlap, i.e., propositions having identical concepts or anaphoric ties. This process occurs at no cost to the limited-capactiy system. There is a cost if a matching proposition is not found in short-term memory. In that case a search of the text base in long-term memory is initiated. Thus, while I speak of assembling and integrating propositions requiring capacity, this does not exactly reflect the assumptions of Kintsch and van Dijk.
3. For example, when asked how much money you have, you would be *truthful* to say "I have $10" when you have $50. If you have $50 you also have $10. But there is a strong constraint on pragmatic grounds (see Grice, 1975), against imprecision. This constraint applies as well to the case here.
4. One difference is whether a model assumes a hierarchical structure above the level of event chains. It is clear that the concept of levels of importance must be captured by some model of the text. There could of course be a level of text model that essentially summarizes event chains, thus

constructing another level. However, the kind of model suggested by Figure 3-4 would directly represent the concept of focal importance.

5. The information in sentence (5) that the scene was in a part of Manhattan where one would ordinarily expect air conditioning is a potentially important clue to future developments. In fact, the reader may immediately sense that a mystery is developing. How *soon* the reader updates the text model will depend on such things. Of course authors of mystery and detective stories depend on such anticipatory devices.
6. There is some evidence that when a nonhierarchical event-based description is compared with a story grammar, each independently can account for subjects' recall (Omanson, 1982b). However, it is possible that this would not be the case for an event theory with a more hierarchical structure. There is some evidence that reading times are affected by story grammar categories (Mandler & Goodman, 1982), but the eye-fixation times of the Just and Carpenter (1980) experiments were unaffected. In either case, the issue is whether the story grammars are powerful schemata only to the extent that they capture event sequences.