



MICROCOPY RESOLUTION TEST CHART
 NATIONAL BUREAU OF STANDARDS-1963-A

DOCUMENT RESUME

ED 137 753

CS 003 367

AUTHOR Ortony, Andrew
 TITLE Remembering and Understanding Jaberwocky and Small-Talk. Technical Report No. 28.
 INSTITUTION Bolt, Beranek and Newman, Inc., Cambridge, Mass.; Illinois Univ., Urbana. Center for the Study of Reading.
 SPONS AGENCY National Inst. of Education (DHEW), Washington, D. C.
 PUB DATE Mar 77
 CONTRACT 400-76-0116
 GRANT HEW-NIE-G-74-0007
 NOTE 36p.

EDRS PRICE MF-\$0.83 HC-\$2.06 Plus Postage.
 DESCRIPTORS *Cognitive Processes; *Memory; Models; Psycholinguistics; Psychological Studies; *Reading Comprehension; *Reading Research; Recall (Psychological); Retention; *Semantics
 IDENTIFIERS *Center for the Study of Reading (Illinois)

ABSTRACT

Since not everything that is understood is remembered and not everything that is remembered is understood, models of language processing should be able to make a distinction between comprehension and memory. To this end, a case is made for a spreading activation process as being the essential ingredient of the comprehension process. Concepts activated during comprehension restrict the search set for candidate concepts and also constitute part of an episodic representation that can become part of long-term memory. The way in which these representations atrophy is discussed, as is the way in which their idiosyncratic components are eliminated in producing representations in semantic memory. Some observations on the comprehension and memory of text are made, and arguments are presented to show how intrusions and omissions in recall can be handled. Existing experimental data are reanalyzed in terms of the proposed model and alternative interpretations consistent with the model are shown to be possible. (Author/AA)

 * Documents acquired by ERIC include many informal unpublished *
 * materials not available from other sources. ERIC makes every effort *
 * to obtain the best copy available. Nevertheless, items of marginal *
 * reproducibility are often encountered and this affects the quality *
 * of the microfiche and hardcopy reproductions ERIC makes available *
 * via the ERIC Document Reproduction Service (EDRS). EDRS is not *
 * responsible for the quality of the original document. Reproductions *
 * supplied by EDRS are the best that can be made from the original. *



CENTER FOR THE STUDY OF READING*

Technical Report No. 28

REMEMBERING AND UNDERSTANDING JABERWOCKY
AND SMALL-TALK

Andrew Ortony

University of Illinois at Urbana-Champaign

March 1977

University of Illinois
at Urbana-Champaign
1005 West Nevada Street
Urbana, Illinois 61801

Bolt Beranek and Newman Inc.
50 Moulton Street
Cambridge, Massachusetts 02138

This work was supported in part by a Spencer Fellowship awarded by the National Academy of Education, and in part by the National Institute of Education under Grant No. HEW-NIE-G-74-0007 and Contract No. MS-NIE-C-400-76-0116.

*Designation as a Center pending approval.

MS 003 367

Remembering and Understanding Jaberwocky and Small-Talk

If a model of human cognition is really to be a model of human cognition, one of the things it ought to do is to distinguish human cognitive capacities that are in reality distinct. Two capacities that current models in both psychology and artificial intelligence (AI) generally fail to distinguish is the capacity to understand and the capacity to remember. The usual theoretical strategy that existing models employ seems to be to convert an input string into an underlying representation of its meaning, and to store that representation in memory as an addition to the knowledge base. Now this is not all wrong, but nor is it quite right. People understand things they don't properly remember, small talk, for example; and they remember things they don't properly understand, Jaberwocky being a case in point. By failing to distinguish comprehension and memory, we may fail to capture essential aspects of each, and the models we adduce may be correspondingly inadequate. Accordingly, one of the major issues to be addressed in this paper is the difference between comprehension and memory.

A second main thrust in what follows will concern the processing of text. Traditionally, emphasis has been placed on the processing of fairly small units such as the sentence. To be sure, some theoreticians have considered larger structures such as paragraphs or even short stories, but for the most part even the more ambitious attempts at dealing with larger, structured, passages have been rather restricted in scope and have often amounted to little more than treating them as sequences of individual sentences. One phenomenon of human language processing that should be

of concern to the cognitive scientist is the distortion and omission in recall of material that was processed and understood. People are not tape recorders, or video tape recorders. A model of language comprehension which stores a representation of the meaning of an input and then regurgitates it in toto is not a model of human language comprehension, although it might be a model of superhuman language comprehension. Our models need to reflect what humans actually do do, not what superhumans might do.

At least within the computer simulation community, the dominating concept in natural language comprehension work today is the concept of a "frame" (Minsky, 1975; Charniak, 1975; Winograd, 1975). Other researchers have worked on related, although sometimes rather different notions such as "scripts" (Schank & Abelson, 1975; Schank et al., 1975), and "schemata" (Bobrow & Norman, 1975; Norman, 1975; Rumelhart, 1975; and Rumelhart & Ortony, 1975).

Frames, or "schemata" as we shall call them, are interacting structural representations of domains of knowledge. Their purpose is to represent the general knowledge required to comprehend language and to provide the bases for interpreting incoming information. The idea is that incoming information will suggest schemata in a predominantly bottom-up fashion until good candidates are established. While this happens, top-down processes contribute to the further selection of schemata for subsequent processing. An important feature of schemata is that they permit easy inferences to be made about certain unstated aspects of an input. Thus, for example, if (1) is encountered

(1) John went to get some things at the supermarket

the "supermarket" schema will contain information about baskets, thus enabling the inference that John put his purchases in a basket to be made. In linguistics this bringing to the fore of relevant selected information has been called "foregrounding" and it is discussed at length by Chafe (1972).

There is no doubt that the general notion of schemata is a useful one. It clearly provides a representation of knowledge which permits the collection of pertinent relationships and concepts together, but theorists have tended to regard schemata as a panacea, often treating them as an alternative to other representational suggestions such as semantic networks or demons. To some extent this latter claim is an exaggeration. Winograd (1974), for example, suggests the possibility that such a representational system could be combined with others to provide the necessary richness. But few seem to have paid much attention to this kind of proposal.

It is one of the purposes of this paper to suggest that the emphasis has been in the wrong direction. Schema-based conceptions of the organization of memory concentrate too much on the nature of concepts in memory, rather than on the nature of memory overall. We will argue that the best candidate for representing the overall structure is still the semantic network, and the best candidate process to operate on it is spreading activation.

In their recent modification and improvement of the Quillian (1968) network model of semantic memory, Collins and Loftus (1975) introduce some additional processing assumptions. The first is that when a concept is processed activation spreads from it in a decreasing gradient; the second is that release of activation from a concept continues at least as long

as that concept is processed; and a third relates to decrease of activation over time. The fourth addition is that activation from different sources summate and that there is a threshold which determines whether or not an intersection is found. Added to these are two additional structural features. First, that semantic similarity plays a larger role in the organization of the network, and second that the names of concepts, i.e. words, are stored in a lexical network which is to some extent independently "primeable."

Suppose we change the first assumption to a "restraining" principle, namely, that if a concept does not receive sufficient activation to exceed its threshold, then it does not transmit activation to neighboring concepts. Suppose further that each node is a unit name for a concept which is itself a schema with an internal structure. We would then have two levels of associations. The first level would be between-concept connections, a network representing memory as a whole. The second level would be within-concept connections, representing the conceptual relations inherent in individual concepts or schemata.

The between-concept connections are those traditionally represented in semantic networks. For example, as Collins and Loftus represent it, part of the network might contain connections between the nodes for, say, "fire-engine" and those for "ambulance," "car," "truck," "bus," "fire," "red," etc. By contrast, the within-concept connections are rather different. Thus, for example, the schema for "fire-engine" would probably not contain references to all the concepts related to it in the semantic network. Rather it would be a mini-network representing general knowledge about

fire-engines. It would represent the knowledge that they were fast-moving road vehicles, that they carried equipment and people for fire-fighting, that they carried ladders and hoses, and that they make a great deal of noise, etc. The within-concept, schema-based representations are thus much more highly structured, not representing more-or-less arbitrary connections between concepts, but representing the actual relationships between those concepts required to represent knowledge about the particular domain.

With this kind of representation, there are two distinct ways of getting from one concept to a token representation of another. One can move between concepts by simple association, thus, for example, getting from "fire-engine" to "ambulance." Or one can move within a concept and get from "fire-engine" to, say, "hose." In some cases, it is reasonable to suppose that one can get to a concept by both between-concept associations and within-concept associations, "fire-engine" to "red" might be an example.

Since schema theorists generally suppose that individual schemata can be regarded as the nodes in a network, the suggestions above may appear to add little to what is already a fairly standard approach. There are, however, some important differences. One is that some of the connecting links between concepts will be simple association links, while others are the labeled relations that participate in a particular schema. Some, but not all, of the labeled links are also simple associative links. But, more important than this, is the fact that spreading activation need not operate upon the whole network (i.e. both within and between concepts) but in the first instance need operate only at the between-concept level.

The use of spreading activation will provide us with a basis for distinguishing comprehension and memory, particularly if this notion is coupled with a distinction between episodic and semantic memory.

Comprehension and Memory

In order to distinguish comprehension and memory one of the first tools we will require is the episodic/semantic memory distinction. The distinction between episodic and semantic memory as envisaged by Tulving (1972) is a pretheoretical one, one to assist in theory construction. He indicates certain characteristics that each might have: episodic memory is autobiographical, it contains spatiotemporal information about experienced events, and it incorporates perceptual information; semantic memory is a mental thesaurus. Episodic memory, in its representation of experienced episodes, represents these episodes in an idiosyncratic way. It does so, as it were, in a way that is close to the way in which the stimulus was subjectively experienced. This will become more clear as we proceed. While a distinction between episodic and semantic memory is now fairly well accepted in psychology, it has met with some resistance in Artificial Intelligence (AI) circles; Schank (1975) has explicitly argued against it.

It is now widely acknowledged that comprehension involves the utilization of knowledge. Memory, on the other hand, involves the retrieval, recognition, and in some cases even the regeneration of representations of knowledge. In both comprehension and memory, the knowledge involved can be either semantic or episodic. Suppose that when we read a sentence (or

perceive an event), bottom-up processes initiate a spread of activation in the semantic network. There is presumably a context in which the input is encountered and this context will itself give rise to activation of a set of concepts. So, the context together with the input will set up a pattern of activation, spreading from all those nodes whose thresholds are exceeded. The restraining principle suggested earlier, contributes to the prevention of an indiscriminate spread to every node in the network. Given no further input, the number of activated nodes will reach a maximum and start to decline. The subgraph formed by that subset of nodes activated, we call the final activated subgraph, for it represents the subgraph of the semantic network which, in containing the greatest number of activated concepts, represents the best final list of concepts involved in processing the input. Each concept in the final activated subgraph has an excitation level greater than or equal to its threshold, and each permits access to its corresponding schema for deeper processing and inference.

The spread of activation through the network is one of the crucial components in comprehension. It serves to restrict the size and content of the search set for the top-down employment of schemata. It is a context-sensitive process, since it produces a different final-activated subgraph for the same input under different contextual conditions, and it permits of different levels of processing (see Craik & Lockhart, 1972) insofar as the size of the final-activated subgraph can be "artificially" varied. Different levels of processing can also be attained by engaging in more or less inference, which is achieved to a great extent by processing at

the level of schemata rather than at the between-concept, semantic network, level. Taking this approach, the process of comprehension involves activating concepts related to those of the input and the context, and engaging in inferences based primarily on those concepts. Put this way, it becomes clear that comprehension is a process prior to, and distinct from, the creation of a memorial representation.

The generation of an episodic representation in memory clearly has to take advantage of the concepts employed in the comprehension process, and in order to do this, it would appear necessary to propose episodic representations that are rather more complex in structure than past research has been willing to allow. In particular, this suggests a tripartite organization of episodic representations, an organization which includes surface structure, semantic representation, and input associates. The proposal is that an episodic representation is generated as a concept centered around the theme or topic. The representation would include the final activated subgraph by including tokens of all its concepts. These tokens would allow access between the episodic representation and knowledge in semantic memory which is either related to the episodic representation itself, or to its components. It would get built up in short term memory and then entered into episodic memory as a record of the experienced meaning of the input.

Such a proposal for tripartite episodic representations has a number of interesting features, particularly if a decay function is associated with each component of the representation. One such feature concerns the relationship between the surface structure and the semantic representation.

~~Details of the semantic representation which were not explicit in the~~
surface structure are normally determined with the help of concepts in the final activated subgraph, and this in turn depends for its particular character on the context in which the input was encountered. Thus, if the surface structure is temporarily stored but inadequately processed, as might happen when one hears but doesn't "register" what someone says, a reinterpretation based on the surface structure alone may give rise to a slightly different episodic representation. Thus, the semantic representation, which can be thought of as an instantiated schema, may have different values for some of the variables whose values were inferred.

Whereas it can happen that occasionally a representation of surface structure survives without the other two parts which constitute an episodic representation, normally the situation is the other way round. That is, the decay functions associated with each level of representation are such that the surface structure is lost first, the semantic representation second, and the final activated subgraph last. At each decaying level, the constituents may themselves be decaying at different rates, so that not all constituents of all levels are equally available at different times. Each level, however, provides some reconstructive potential for the more rapidly decaying level next to it. Thus, if the surface structure has decayed the semantic representation provides the possibility of reconstructing it in whole or in part, although more often than not, this reconstruction results in a paraphrase rather than a verbatim recall. More interesting, however, is what happens if details have been lost from the second level,

the semantic representation. If, for example, the value assigned to a particular variable has decayed, the bottom level provides a list of candidate concepts for regenerating or reconstructing the semantic representation. Later, it will transpire that this is the way in which one can account for errors, omissions and intrusions in recall. This reconstructive potential, however, is not perfect since the contextual details, while influencing the final activated subgraph, are not explicitly stored in the top level of surface structure representation. So, if one can remember a speaker's exact words (surface structure), one can generally remember (or at least reconstruct) how it was interpreted and, consequently, at least some of the associated concepts. However, one may have forgotten the surface structure and still have a representation of the semantic structure, or even lost that but still have some associated concepts left. The associations represented by the final activated subgraph must be represented because they are in part determined by the total context in which the input originally occurred; this is how encoding specificity can be handled (Tulving & Thomson, 1973). Notice that so far nothing has gone into semantic memory; semantic memory has merely been invoked in getting episodic representations built up.

On the view being proposed, the knowledge structures stored in semantic memory are for practical purposes relatively permanent; over time they may undergo changes, but semantic memory is basically a repository for encyclopedic knowledge. Episodic memory, on the other hand, is memory for the kind of information one might incorporate in a personal diary (cf. Ortony, 1975).

However, it is more than a record of personal experiences. It is also the source of knowledge for semantic memory, provided that the necessary processing takes place before too much of the episodic representation is lost. The contents of episodic memory impose certain restrictions on the way that knowledge can be represented. For the sake of simplicity, let us, for the moment, restrict our discussion to knowledge represented propositionally.

Semantic memory represents knowledge that a person believes to be true. While it is derived from experience, it is a record of what he knows independently of experience. Episodic memory contains this knowledge embedded in source and modality information but it cannot be inferred directly from it. A simple example: I read a book that claimed that X might be represented as part of the contents of one entry in episodic memory. X here represents some arbitrary proposition. A normal person satisfied that he wasn't deluded would probably easily move to A book claimed that X, and that might or might not be entered into semantic memory depending on the relevance of the book to what is already known. But the putative knowledge is X, stripped of source and circumstances, and a complicated, time-consuming process of evaluation may be required before a person would be willing to assent to X alone. If time is available, and relevance and potential use in inference seem to justify it, the evaluation may be undertaken and, if X can stand alone, it may be transferred to semantic memory alone. The evaluation required to get from A book claimed that X to X is a form of inference. So, we have a situation in which the

comprehension of an input item depends on accessing information in semantic memory so as to enable the construction of a representation in episodic memory. The putative knowledge in episodic memory (the X in A book claimed that X) can subsequently be freed for semantic memory if it is sufficiently relevant and if sufficient subjective truth can be accorded to it. These evaluative processes, of course, need not be conscious.

A Re-interpretation of Some Data

An associative model of the kind just outlined appears to be capable of distinguishing between comprehension effects and memory effects. Although current fashion in cognitive psychology tends to be rather unsympathetic towards associative models, it may be that associationism in general has been confused with particular associative theories, or pieces of them (e.g., Anderson & Bower, 1973; Collins & Quillian, 1969). If this is indeed true, then it is necessary to show that some of the experimental results often cited as good reasons for rejecting associative accounts need to be re-examined. The experiments I have chosen to look at are two that were conducted in our own laboratory.

Ortony and Anderson (1977) constructed a number of sentence sets in which a noun phrase, either a name or a definite description, was paired with one of two predicates. For example, one of the sentence sets that was used was:

- (2) ~~The inventor of dynamite had a profound effect on the nature of warfare...~~
- (3) Alfred Nobel had a profound effect on the nature of warfare...
- (4) The inventor of dynamite wore a fine beard...
- (5) Alfred Nobel wore a fine beard...

Each subject in the experiment saw two sentences from each set. In the example above, a subject would see either (2) and (5), or (3) and (4), and he was then asked to write a continuation of the sentence with the word "but." The purpose of the task was to encourage the subject to think carefully about the meaning of the sentence. After the subject had worked through the list he was given a recognition task in which all the sentences from all sentence sets were shown to him. His task was to determine for each sentence whether it was a new sentence (one which he had not seen earlier) or an old one (which he had). The primary purpose of the experiment was to investigate the relationship between names and descriptions, and between them and different kinds of predicates. This was done by analyzing the recognition errors subjects made. Our prediction was that if a subject saw a pair such as (2) and (5) he would be less likely to falsely recognize (3) and (4) since the first, (2), is about the inventor of dynamite and can be judged and understood irrespective of who he was, and the second, (5), concerns the personal characteristics of a particular man. This we felt was not true with sentences (3) and (4), where, it was felt, a subject might have implicitly substituted, or at least have accessed, the name of the described individual in (4), and a suitable description of the named individual in (3). These predictions were confirmed.

One of the implications of this experiment, as with Anderson and Ortony (1975) to be discussed shortly, is that the noun phrases in subject position in these sentences cannot be appropriately interpreted until the predicate that follows them has been processed. The context that the predicate supplies,

guides and constrains the interpretation of the noun phrase. This result, it seems, is one that can be well handled by a spreading activation model of the kind proposed in this paper.

Assuming only that semantic memory contains, or can generate the concept of the inventor of dynamite, when it is encountered, activation will begin to spread to adjacent nodes, not only to inventor and dynamite, but perhaps also to explosion, gunpowder, war, genius, scientist and so on. While this activation is taking place the predicate begins to get processed. If the predicate is had a profound effect on the nature of warfare, we might expect activation from the predicate, or concepts within it, of such concepts as war, guns, death, etc., but not of genius and scientist; the result would be that concepts in the war domain would have their activation increased while there would be no support for the genius, scientist concepts whose activation levels might thus be expected to fall off. Hence the probability of creating an activated subgraph containing war concepts would be much higher than that of a subgraph comprising concepts about smart scientists. But, if the second predicate is encountered, the predicate wore a fine beard, concepts in the domain of people would be activated--concepts perhaps including man, scientist, genius, eccentric and so on. In this case we have activated both war concepts and people concepts and the resulting activated subgraph is more in doubt--it might contain war concepts only, people concepts only, or both. In any case, there will clearly be more activation of the war concepts given the war predicate, than there will be given the beard predicate because the latter can contribute virtually

nothing. Whether or not the subject knows the identity of the inventor of dynamite is not a critical question except to the extent that if he doesn't, there is no chance that it can get activated at any stage during the comprehension process. If he does know it, and even if it does get activated during the process, it doesn't make any difference unless it remains activated at the end of the process and is part of the final activated subgraph. This brings us to the question of the strength of the association. It is presumably the case that George Washington is much more likely to be activated by The first President of the United States than is Christian Barnard by The first man to do a heart transplant. Nevertheless, which nodes fall below threshold will also depend on the degree of connectivity in semantic memory. Consequently The inventor of dynamite might be just as stable (i.e., unlikely to fall below threshold) when related to the concepts derived from the war predicate as Alfred Nobel would be, even if Alfred Nobel were very strongly associated with the conceptual structure deriving from the description. So, a network model could explain these results in terms of differing combinations of the various activated concepts. Notice that we use the notion of "intersection" but not of "intersection search." An "intersection" is merely a concept activated from two or more sources. Its effect is to boost that concept's level of activation and consequently the probability that it exceeds threshold and thus activates concepts related to it.

Anderson and Ortony (1975) had subjects learn two sentences from each of fourteen sentence sets. For example, one of the sentence sets used was:

- (6) Nurses are often beautiful.
- (7) Nurses have to be licensed.
- (8) Landscapes are often beautiful.
- (9) Taverns have to be licensed.

Each subject was asked to learn two from each set: in this example, a subject would have learned either (6) and (9), or (7) and (8). If he had seen this latter pair, then of the two retrieval cues ACTRESS and DOCTOR, which he later received in a memory test, DOCTOR was regarded as the close cue and ACTRESS as the remote cue for his target sentence Nurses have to be licensed. The landscape sentence acted as a control. What was found was a pattern such that DOCTOR would be a superior cue for Nurses have to be licensed, but that ACTRESS would be superior for Nurses are often beautiful. The cues were not effective by virtue of their association with either the subject alone or the predicate alone. Doctor was not effective by virtue of its connection with licensing because it never elicited the control sentence, Taverns have to be licensed, and had its effectiveness been due to its connection with Nurse then it should have been equally good at eliciting both of the nurse sentences. The conclusion, therefore, was that the effectiveness of the close cue was due to its relation to a meaning for the entire sentence, rather than any of the individual parts, alone.

These results can be accounted for in exactly the same way as for Ortony and Anderson (1977). If both activated, Nurses and License are, going to activate health professional and doctor with a greater probability and to a greater degree than are nurse and beautiful, so that doctor is

more likely to be a concept in the final activated subgraph than is actress. Actress on the other hand, is closer to woman which is perhaps a better candidate for final activation than doctor starting from nurse and beautiful.

A retrieval cue is no more than a stimulus which is semantically and/or experientially related to a record of an earlier stimulus. If we could assume that the final activated subgraph stayed in memory for some time without losing its integrity, then all a cue would need to do would be to reactivate or make contact with some part of the subgraph. Consider the word ACTRESS in its relation to nurses are often beautiful. There are two possibilities: the concept either does or does not appear in the final activated subgraph. If it does not, we can be fairly sure that woman does because it will have received activation from both subject and predicate. In this case, the cue ACTRESS will activate the actress concept which in turn will lead to activation on neighboring nodes such as woman. If the activation of woman is sufficiently increased by actress, then the entire subgraph could increase its overall level of activation giving rise to subsequent retrieval of the subgraph. If the subgraph was not dense the additional activation of woman might give only partial retrieval of that subgraph. It is clear that the same kind of process would lead to recall in the case where the cue was itself a concept which was in the final activated subgraph, or where the cue was any non-linguistic input (such as a picture) which made eventual contact through its concept(s) with a concept in the original.

How do episodic and semantic memory combine to enable a subject using a retrieval cue such as ACTRESS to retrieve from (episodic) memory Nurses

are often beautiful more reliably than Nurses have to be licensed? The cue activates the concept actress in semantic memory which in turn activates, let us say, woman and beautiful. The subject, being told that he has to recall the recently learned sentences, has other information available: source and circumstance information which together with the activated concepts in semantic memory, should allow access to a representation in episodic memory produced during the acquisition stage of the experiment. The subject then selects that representation in episodic memory with the most of his currently activated concepts represented in it. There are several ways in which he can go wrong. First, some of the associated concepts in the episodic representation might have atrophied. Second, the activated concepts may not be unique to one episodic memory representation, in which case the subject will retrieve the correct response only some of the time. Thus, we can account for the kinds of responses a subject might make. We can even account for guesses and no response at all. In the first case there could be too few of the associates in the episodic memory representation. In the second case there could be none, due to inadequate processing at comprehension time. In both cases no unique episodic representation would be accessible.

Representing Text

So far we have pretty much restricted our account to the comprehension and memory of individual sentences. But most language processing that people do involves linguistic units of much greater length and complexity than individual sentences. Whereas the sentence is a fairly standard unit

of linguistic analysis, it probably does not have any particularly important psychological status as such. With inputs sufficiently short to enable the construction of a well-integrated final activated subgraph the entire message is represented. On the other hand, as the unit of input becomes longer, there would probably be a problem of maintaining activation of all the relevant concepts for a sufficient length of time. If a reader does not attempt to break such a message down into more manageable "chunks," corresponding to smaller, better integrated subgraphs, the chances are that some of the important concepts involved in comprehension will be lost. The construction of several related episodic representations to help in the representation of a complex or a long message leads to the idea that such messages will have to be represented in higher level structures. We might consider how this could work with the individual units combining to give a representation of a passage of text. One possibility is through some kind of chaining process. Such a process would allow an event-based representation of the total meaning of the passage to be built up in episodic memory.

Each input unit (sentence or whatever) produces, as we have seen, a final activated subgraph in semantic memory. The representation in episodic memory incorporates tokens of, or pointers to, these activated concepts. The meaning of the entire passage would be cumulatively constructed by generating connected episodic representations of psychologically discrete components. Each unit that is processed produces a pattern of activation which is in part determined by the pattern resulting from the last processed unit. If there is a dramatic change of subject, it is to be expected that

the existing pattern of activation will not produce many intersections so that recently activated nodes will rapidly lose their activation while a new set becomes activated. A further feature is that the final activated subgraph critically depends on each constituent (and its order) which went into its making. At the same time it need contain no distinguishable

representation of each of those constituents. It has been built up through an interaction between the input and the effect of the context (including previously processed input) on semantic memory. It at once contains more and less than the original passage, a desirable feature since at an anecdotal level, at least, it accords with our daily experiences in reading text. At the same time, the semantic structure part of the representation will grow when the topic is the same, but will not provide suitable points of attachment for radically new material.

In reading a passage, the reader picks up clues as to the structure of the input from his knowledge of syntax aided by the punctuation. In listening, these clues are provided by prosodic features of the utterance. Without such clues, a language processor has no basis for chunking the input into psychologically discrete units for processing. We can assume, therefore, that in normal circumstances, a human language processor can identify units to be processed as units, at various levels of structure. As he reads a passage, each unit provides a potential semantic structure into which the next unit can be absorbed, and each unit provides a pattern of activation (or a context) for the current unit to influence. Since people do not remember passages word for word, we have to assume that the surface

structure of each unit is not normally stored. Further, by building onto existing semantic representations, the overall order in which the original units were processed may not necessarily be represented. Sometimes, when order is important, special strategies may result in the establishment of independent (but selected) episodic representations. Such a model, one in which there is no distinct representation for the meaning of individual constituent inputs outside of the larger context (in this case the passage) in which they occur, accords well with the findings of Bransford and Franks (1971) and Spiro (1975) which indicate that subjects do incorporate incoming information into a whole in such a way that meaning components can be extracted but which makes it difficult to reliably isolate individual inputs.

When we read a sentence what we "get out of it" is not a simple function of the concepts whose tokens occurred in it. We get a complicated representation which may to a greater or lesser extent involve all kinds of additional associated concepts. Yet when we think about or talk about that sentence, we seem to be able to operate on a more condensed comprehensive representation. The same is true of the relationship that these condensed versions bear to the representation of a larger block of text. Our more superficial, least detailed representation of a paragraph or a novel does not contain all the concepts which all the tokens in it activated; indeed, it does not contain all the condensed representations of its individual units. Something new and more succinct is created. I have a reasonably good representation of Crime and Punishment, an entire novel, but it is not a

conglomerate of all the concepts and their associates that were activated while reading it. My representation is too vague and skimpy for that. I know that it was about a man (Raskolnikov) and his actions (a hatchet job on an innocent old lady) and his subsequent emotions in dealing with a police inspector who knew Raskolnikov had committed the crime but who couldn't prove it.

On the model here being proposed, a novel would be represented as a connected cluster of episodic representations, perhaps all connected to a node for it, say, "Crime and Punishment." The number of episodic representations need not be very large, especially after a long time. These constituent episodes each contain some remnants of their final activated subgraphs, but many of the values for particular variables within them may have been lost (for example, the name of the inspector, or the old lady). The overall representation of the novel also contains its own final activated subgraph.

Both at the level of sentences, and at that of text, the kind of episodic representations that have been proposed permit of a reasonable account of distortions and omissions in recall. As was suggested earlier, the levels within an episodic representation allow communication between them. We further suggested that a rate of decay is associated with each level, very fast for the top, surface structure level, and less fast for the lower levels. At the lower levels we probably need to think in terms of a differential decay rate for the individual concepts involved. If I tell you today that my dog is sick, and two weeks later you politely inquire

as to the health of my ailing cat, we need to be able to account for your misconception. The structure we have attributed to episodic representations permits us to deal with this kind of phenomenon. It is necessary to assume that the semantic structure component of the episodic representation has in some way atrophied. In the particular example cited, we have to assume that what has been lost is the representation of what was the particular creature that was sick. What remains is a structure representing the information that something or other belonging to the person who made the original utterance was sick. In recall it is therefore necessary to reconstruct a representation of the original. The record (or what remains of it) of the final activated subgraph which constitutes part of the episodic representation provides a set of candidate concepts for engaging in this reconstructive process. Clearly, these candidate concepts will have to be selected, presumably on the basis of inferential processes. It is beyond the scope of this paper to speculate as to how these processes operate. The main purpose of introducing the issue is to suggest a way of accounting for intrusions in recall. The utilization of a decaying representation of a final activated subgraph affords the possibility of producing recalls which utilize concepts different from, but related to, those which appeared in the original.

Conclusion

We have seen how spreading activation operating on a semantic network can be used to distinguish what is understood from an input from what is remembered of it. In doing so, we have postulated frame-like entries, which

we have called schemata as the basic units in the network. An attractive addition would be to allow the spreading activation to also extend to representations in episodic memory, thus allowing entire experienced events to play the role of associates to inputs to and concepts in semantic memory. In this way episodic representations, or at least tokens of them, would be able to feature as associates of later ones, thereby approximating more nearly to what we know happens to people, namely that episodic representations can relate to and be suggested by, not only general concepts (semantic memory), but also by one another. Consequently, we compensated for the relative informational poverty of semantic network theories while retaining their power to selectively activate concepts related to the input taking into account the context in which it is encountered. Concepts in memory are not to be regarded as structureless points but as richly structured representations. Even concepts which have been traditionally regarded as simple, unanalyzable properties, such as those designated by color terms, have an internal structure. Halff, Ortony, and Anderson (1976) conducted a study investigating the word "red." One of the conclusions of the study was that the internal representation of a concept such as red has to be a real interval and cannot be a point, as implied by current models of semantic memory.

It might seem philosophically heretical to suggest that there are no simple, unanalyzable properties, since, traditionally, in western thought there have been, and color terms were frequently taken as paradigmatic examples. Yet, we lose little if we relinquish this dogma. Concepts may

be more or less structurally complex without it being necessary that there be some ultimate, simple (therefore presumably structureless) concepts. Perhaps we should distinguish not between simple and complex, but between complex and compound. The concept airplane we might say is compound in that it incorporates other concepts; it is a compound complex concept. The concept of red or of warm might be regarded as a complex concept, but not a compound one. The structure of the concept might perhaps be regarded as some kind of distribution (see Anderson, 1975). Certainly it is straining the conventional sense of "simple" to call such concepts simple. They may be relatively simple, they may even be unanalyzable, but certainly they are not absolutely simple, structurally empty "things." Simple ideas may be useful for philosophers, but are not, I suspect, for psychologists.

In spite of the bad reputation that associationism has, "pure" associationism is innocuous; indeed, in a sense it is impossible to reject for it asserts only that ideas are related in a non-arbitrary way. To reject such a claim is senseless and leaves one without even the basis for a theory. However, by making certain assumptions one can in fact explain a number of phenomena which critics of associationism have considered beyond the domain of associative model. It is certainly not appropriate to try to explain all of cognition in this way. Associationism constitutes a good model only of certain aspects of memory and understanding; but that is only part of the story. Long-term memory comprises concepts of varying degrees of complexity, related to one another by directed associative connections of various types. The totality of inputs to this network has various

short term and longer term effects. The short term effects include the excitation of a subgraph of related concepts, the long term effects include the generation of new concepts, new connections between concepts, and changes in associative strengths, i.e. relative accessibility between concepts.

A possible objection to such a system might be that it is too passive. But this is not a serious problem. Firstly, in it meanings are constructed--that is, representations in episodic memory are created by reference to semantic memory and sometimes recreated or reconstructed for recall. It might be objected that all the relations in semantic memory are given and static. This is partly true, but it is not objectionable, for the detailed structure of memory is forever changing, both with respect to content and with respect to strengths of associative links. Further, the relationships that are relevant between concepts will be largely determined by the contribution that context makes to activation levels in semantic memory. It might further be objected that the system is too S-R bound--that all that happens is an inevitable result of the activation of some concept or concepts by an input. It would then be argued that this leaves no room for a hypothesis/test strategy in perception or comprehension. The answer to these objections could be along the following lines. Whereas it is true that some of what happens is a direct result of activation by an input, it does not follow that it all is. If every concept has a threshold it is reasonable to suppose that some are lower than others. It is indeed the case that sometimes all that is needed is that the input activates a concept--such activation may even arrest other processing. Bobrow and

Norman (1975) consider the case of people so occupied with a task that they cannot "hear" words directed to them. Yet, frequently, in these situations, people do respond if their name is spoken, or if the word fits into the context of their absorbing activity. Our model readily accounts for this--one's own name is connected to a low threshold concept (of self); words which fit the context provide additional activation to concepts already receiving some--so less, if any, is required to reach threshold. These cases of "perceiving without hearing" are precisely cases in which a certain amount of passivity is entirely appropriate in an explanation. The subject has no control over his sudden perception, it happens to him, he doesn't do it. But, there are some things that are not so simple. People can and do exercise some choice over what they are thinking about. Given that a person may have certain goals, he may (rightly or wrongly) decide that he is more likely to achieve them by concentrating on one group of concepts rather than another. Such a choice may give rise to activation of one area of his network rather than another. Thinking may be partly the selective increasing of activation in semantic memory or at least the temporary changing of thresholds. Passivity obviously is not the whole story. Concepts may embody procedures or hypotheses; explorations in memory may test them.

Inevitably, more processes would have to be specified to properly characterize all the phenomena with which such a theory should be concerned. Nevertheless, perhaps that old associationist theory of memory has life in it yet. Ironically, its vagueness is at once its power and its weakness;

its power because it can be shaped almost at will; its weakness because unshaped or badly shaped it is vulnerable to criticism from every quarter. But at least insofar as it seems capable of permitting a distinction between memory and comprehension, it has much to commend it.

References

- Anderson, J. R., & Bower, G. H. Human associative memory. Washington, D. C.: Winston & Sons, 1973.
- Anderson, R. C. Knowledge of variability and the contextual determination of meaning. Unpublished manuscript, University of Illinois at Urbana-Champaign, 1975.
- Anderson, R. C., & Ortony, A. On putting apples into bottles: A problem of polysemy. Cognitive Psychology, 1975, 7, 167-180.
- Bobrow, D. G., & Norman, D. A. Some principles of memory schemata. To appear in D. G. Bobrow and A. M. Collins (Eds.), Representation and understanding. New York: Academic Press, 1975.
- Bransford, J. D., & Franks, J. J. The abstraction of linguistic ideas. Cognitive Psychology, 1971, 2, 331-350.
- Chafe, W. Discourse structure and human knowledge. In J. B. Carroll and R. O. Freedle (Eds.), Language comprehension and the acquisition of knowledge. New York: V. H. Winston & Sons, Inc., 1972.
- Charniak, E. Organization and inference in a frame-like system of common sense knowledge. In Proceedings of Theoretical Issues in Natural Language Processing: An interdisciplinary Workshop. Cambridge, Mass.: Bolt Beranek and Newman Inc., 1975. Pp. 42-51.
- Collins, A. M., & Loftus, E. F. A spreading activation theory of spreading activation. Psychological Review, 1975, 82, 407-428.
- Collins, A. M., & Quillian, M. R. Retrieval time from semantic memory. Journal of Verbal Learning and Verbal Behavior, 1969, 8, 240-247.

- Craik, I. M. F., & Lockhart, R. W. Levels of processing: A framework for memory research. Journal of Verbal Learning and Verbal Behavior, 1972, 11, 671-684.
- Halftr, H. M., Ortony, A., & Anderson, R. C. A context-sensitive representation of word meanings. Memory and Cognition, 1976, 4, 378-383.
- Minsky, M. A framework for representing knowledge. In P. H. Winston (Ed.), The psychology of computer vision. New York: McGraw-Hill, 1975.
- Norman, D. A. Resources and schema replace stages of processing. Paper presented at the Sixteenth Annual Meeting, The Psychonomic Society, Denver, Colorado, November 1975.
- Ortony, A. How episodic is semantic memory? In Theoretical Issues in Natural Language Processing: An Interdisciplinary Workshop. Cambridge, Mass.: Massachusetts Institute of Technology, June 1975.
- Ortony, A. & Anderson, R. C. Definite descriptions and semantic memory. Cognitive Science, 1977,
- Quillian, M. R. Semantic memory. In M. Minsky (Ed.), Semantic information processing. Cambridge, Mass.: MIT Press, 1968.
- Rumelhart, D. E. Notes on a schema for stories. In D. G. Bobrow and A. M. Collins (Eds.), Representation and understanding: Studies in cognitive science. New York: Academic Press, 1975.
- Rumelhart, D. E., & Ortony, A. The representation of knowledge in memory. In R. C. Anderson, R. J. Spiro, and W. E. Montague (Eds.), Schooling and the acquisition of knowledge. Hillsdale, N. J.: Lawrence Erlbaum Associates, 1977.

- Schank, R. C. The structure of episodes in memory. In D. G. Bobrow and A. M. Collins (Eds.), Representation and understanding: Studies in cognitive science. New York: Academic Press, 1975.
- Schank, R. C., & Abelson, R. P. Scripts, plans and knowledge. In Advance Papers of the Fourth International Joint Conference on Artificial Intelligence. Tbilisi, Georgia: USSR, 1975. Pp. 151-157.
- Schank, R. C., & The Yale A.I. Project. SAM--A story understander (Res. Rep. 43). New Haven, Connecticut: Yale University, Department of Computer Science, August 1975.
- Spiro, R. J. Inferential reconstruction in memory for connected discourse (Tech. Rep. No. 2). Urbana, Ill.: Center for the Study of Reading, University of Illinois, October 1975.
- Tulving, E. Episodic and semantic memory. In E. Tulving and W. Donaldson (Eds.), Organization of memory. New York: Academic Press, 1972.
- Tulving, E., & Thomson, D. M. Encoding specificity and retrieval processes in episodic memory. Psychological Review, 1973, 80, 352-373.
- Winograd, T. Five lectures on artificial intelligence -- Computer systems for natural language; SHRDLU: A system for dialog; Representation: Formalisms for knowledge; Frames: Some ideas for a new formalism; Conceptual programming: Applying AI to program writing. Edited transcription of five lectures delivered at the Electrotechnical Laboratory in Tokyo, March 18-23, 1974.
- Winograd, T. Frame representations and the declarative-procedural controversy. In D. G. Bobrow and A. M. Collins (Eds.), Representation and understanding: Studies in cognitive science. New York: Academic Press, 1975.

CENTER FOR THE STUDY OF READING

TECHNICAL REPORTS

- No. 1: Halff, H. M. Graphical Evaluation of Hierarchical Clustering Schemes, October 1975.
- No. 2: Spiro, R. J. Inferential Reconstruction in Memory for Connected Discourse, October 1975.
- No. 3: Goetz, E. T. Sentences in Lists and in Connected Discourse, November 1975.
- No. 4: Alessi, S. M., Anderson, T. H., & Biddle, W. B. Hardware and Software Considerations in Computer Based Course Management, November 1975.
- No. 5: Schallert, D. L. Improving Memory for Prose: The Relationship Between Depth of Processing and Context, November 1975.
- No. 6: Anderson, R. C., Goetz, E. T., Pichert, J. W., & Halff, H. M. Two Faces of the Conceptual Peg Hypothesis, January 1976.
- No. 7: Ortony, A. Names, Descriptions, and Pragmatics, February 1976.
- No. 8: Mason, J. M. Questioning the Notion of Independent Processing Stages in Reading, February 1976.
- No. 9: Siegel, M. A. Teacher Behaviors and Curriculum Packages: Implications for Research and Teacher Education, April 1976.
- No. 10: Anderson, R. C., Pichert, J. W., Goetz, E. T., Schallert, D. L., Stevens, K. V., & Trollip, S. R. Instantiation of General Terms, March 1976.
- No. 11: Armbruster, B. B. Learning Principles from Prose: A Cognitive Approach Based on Schema Theory, July 1976.
- No. 12: Anderson, R. C., Reynolds, R. E., Schallert, D. L., & Goetz, E. T. Frameworks for Comprehending Discourse, July 1976.
- No. 13: Rubin, A. D., Bruce, B. C., & Brown, J. S. A Process-Oriented Language for Describing Aspects of Reading Comprehension, November 1976.
- No. 14: Pichert, J. W., & Anderson, R. C. Taking Different Perspectives on a Story, November 1976.
- No. 15: Schwartz, R. M. Strategic Processes in Beginning Reading, November 1976.
- No. 16: Jenkins, J. R., & Pany, D. Curriculum Biases in Reading Achievement Tests, November 1976.
- No. 17: Asher, S. R., Hymel, S., & Wigfield, A. Children's Comprehension of High- and Low-Interest Material and a Comparison of Two Cloze Scoring Methods, November 1976.

- No. 18: Brown, A. L., Smiley, S. S., Day, J. D., Townsend, M. A. R., & Lawton, S. C. Intrusion of a Thematic Idea in Children's Comprehension and Retention of Stories, December 1976.
- No. 19: Kleiman, G. M. The Prelinguistic Cognitive Basis of Children's Communicative Intentions, February 1977.
- No. 20: Kleiman, G. M. The Effect of Previous Context on Reading Individual Words, February 1977.
- No. 21: Kane, J. H., & Anderson, R. C. Depth of Processing and Interference Effects in the Learning and Remembering of Sentences, February 1977.
- No. 22: Brown, A. L., & Campione, J. C. Memory Strategies in Learning: Training Children to Study Strategically, March 1977.
- No. 23: Smiley, S. S., Oakley, D. D., Worthen, D., Campione, J. C., & Brown, A. L. Recall of Thematically Relevant Material by Adolescent Good and Poor Readers as a Function of Written Versus Oral Presentation, March 1977.
- No. 24: Anderson, R. C., Spiro, R. J., & Anderson, M. Schemata as Scaffolding for the Representation of Information in Connected Discourse, March 1977.
- No. 25: Pany, D., & Jenkins, J. R. Learning Word Meanings: A Comparison of Instructional Procedures and Effects on Measures of Reading Comprehension with Learning Disabled Students, March 1977.
- No. 26: Armbruster, B. B., Stevens, R. J., & Rosenshine, B. Analyzing Content Coverage and Emphasis: A Study of Three Curricula and Two Tests, March 1977.
- No. 27: Ortony, A., Reynolds, R. E., & Arter, J. A. Metaphor: Theoretical and Empirical Research, March 1977.
- No. 28: Ortony, A. Remembering and Understanding Jaberwocky and Small-Talk, March 1977.
- No. 29: Schallert, D. L., Kleiman, G. M., & Rubin, A. D. Analysis of Differences Between Oral and Written Language, April 1977.
- No. 30: Goetz, E. T., & Osborn, J. Procedures for Sampling Texts and Tasks in Kindergarten through Eighth Grade, April 1977.
- No. 31: Nash-Webber, B. Anaphora: A Cross-Disciplinary Survey, April 1977.
- No. 32: Adams, M. J., & Collins, A. A Schema-Theoretic View of Reading Comprehension, April 1977.
- No. 33: Huggins, A. W. F. Syntactic Aspects of Reading Comprehension, April 1977.

- No. 34: Bruce, B.C. Plans and Social Actions, April 1977.
- No. 35: Rubin, A.D. Comprehension Processes in Oral and Written Language, April 1977.
- No. 36: Nash-Webber, B. & Reiter, R. Anaphora and Logical Form: On Formal Meaning Representations for Natural Language, April 1977.
- No. 37: Adams, M.J. Failures to Comprehend: A Question of Balance, April 1977.
- No. 38: Woods, W.A. Multiple Theory Formation in High-Level Perception, April 1977.
- No. 39: Nickerson, R.S. & Adams, M.J. Uses of Context in Speech Understanding and Reading, April 1977.
- No. 40: Brown, J.S., & Collins, A. Model-Based Versus Text-Based Reasoning, April 1977.
- No. 41: Pichert, J.W., & Anderson, R.C. Recall of Previously Unrecallable Information Following a Shift in Perspective, April 1977.
- ✓ No. 42: Mason, J., Osborn, J., & Rosenshine, B. A Consideration of Skill Hierarchy Approaches to the Teaching of Reading, April 1977.
- No. 43: Collins, A., Brown, A.L., Morgan, J.L., & Brewer, W.F. The Analysis of Reading Tasks and Texts, April 1977.
- No. 44: McClure, E. Aspects of Code-Switching in the Discourse of Bilingual Mexican-American Children, April 1977.
- No. 45: Schwartz, R.M. Relation of Context Utilization and Orthographic Automaticity in Word Identification, May 1977.