Testing The Psychological Reality of a Representational Model

Dedre Gentner
Bolt Beranek and Newman Inc.

ABSTRACT

A research program is described in which a particular representational format for meaning is tested as broadly as possible. In this format, developed by the LNR research group at The University of California at San Diego, verbs are represented as interconnected sets of subpredicates. These subpredicates may be thought of as the almost inevitable inferences that a listener makes when a verb is used in a sentence. They confer a meaning structure on the sentence in which the verb is used. To be psychologically valid, these representations should capture (at least):

1. Similarity of meaning
   The more similar two verbs seem in meaning to people, the more their representations should overlap.

2. Confusability
   The more confusable two verb meanings are, the more their representations should overlap.

3. Memory for sentences containing the verb
   The sentence structures set up by the verb's meaning should in part determine the way in which sentences are remembered.

4. Semantic integration
   The representations should allow for the integration of information from different sentences into discourse structure

5. Acquisition patterns
   The structural partitions in the representations should correspond to the structures children acquire when they are learning the meanings of the verbs.

6. Patterns of extension
   The representations should be extendible so as to reflect the ways in which people interpret verb meanings when the verbs are used outside their normal context.

7. Reaction times
   The time taken to comprehend a sentence using a given verb should reflect the structural complexity of the verb meaning.

Experiments concerned with predictions 1-5 are described here. The results are promising for a general approach of representation of meaning in terms of interrelated subpredicates, but do not clearly distinguish between several similar representations. For example, to test prediction (2), I read people sentences containing verbs with similar meanings, and asked them to recall the sentences. The degree of overlap in the semantic structures was a good predictor of the number of confusions between sentences. In another sentence-memory experiment (prediction (3)), semantically complex verbs that provided more underlying interconnections between the nouns in a sentence led to better memory for the nouns in the sentence than simple general verbs, or than other complex verbs that did not provide such extra interconnections. To test prediction (5), I tested children's comprehension of a set of possession verbs. Both the order of acquisition among the verbs and the kinds of errors fitted well with an account of the acquisition of verb meaning in terms of interconnected subpredicates.

This research illustrates a breadth-first approach to testing a representation. In the breadth-first approach, many different psychological predictions are made. Each different area of prediction requires a set of process assumptions, and in each case the process assumptions used are those that seem most plausible given previous research in the field. If one representational format can make correct predictions about a number of different kinds of psychological phenomena, then that representation stands a greater chance of being generally useful than one which was tested in only one depth-first way.

This paper describes a program of research that tests a representational format for verb meaning. This research grew out of the LNR (Footnote 1) attempt to the represent the meanings of words in a psychologically satisfying way. Verb meaning seemed a natural place to start for two reasons: (1) verbs are important: it is arguable that they provide the central organizing semantic structures in sentence meanings; and (2) verbs are tractable: their meanings are more easily analyzed than those of, for example, common nouns.

Since different disciplines look at meaning in different ways, it may be worthwhile to describe the stance we took. What we wanted was a system of representation in which we could capture our intuitions about what a word typically conveys; or more specifically about the inferences a person normally makes (or believes should be made) when a word is used. The assumption is that the same representations operate when a person uses the word in speech as when the person comprehends it; however the methodology of experimental psychology makes it natural to spend more time pondering the input process than the output process. This approach differs from thinking of meaning in terms of necessary and sufficient truth-conditions, as many philosophers have done, or from thinking about meaning in generation rather than in comprehension, as many linguists have done. Each of those stances leads to useful intuitions. Overall, there has been a reassuring degree of convergence between the representations proposed.

Representation of Verb Meaning

There are many notational systems for representation of verb meaning (e.g., Abrahamson, 1975; Chafe, 1970; Fillmore, 1971; Gentner, 1975; Lakoff, 1970; McCawley, 1968; Rumelhart & Levin, 1975; Schank, 1972, 1975; Talmy, 1975). These models of verb meaning differ from one another in detail, but there is widespread agreement on the idea that verb meanings can be represented in terms of interrelated sets of subpredicates, such as CAUSE or CHANGE. These subpredicates are not merely concatenated within a word's representation. Rather, they are interrelated, in specific ways. Representations of verb meaning include notation for specifying the relationships among the subpredicates that make up a word's meaning. The notation developed by the LNR Group is a network format. In this system of representation, verb meanings are expressed in terms of subpredicates that stand for states, changes of state, actionals, etc.

The Elements of Verb Meaning. Verbs provide a system in which people can talk about happenings in the world, implicitly distinguishing several types of conceptual possibilities. The simplest of these is the _state_. A stative predicate conveys a relationship that endures for a period of time between two arguments, normally an object (or person) and an object or value within the conceptual field specified by the stative. For example, consider the sentence shown in Figure 1.

Ida owned a Cadillac from 1970 to 1977.

The verb _own_ conveys that a relationship of possession existed between Ida and the Cadillac for some duration. Besides statives for possession, there are a large number of other statives, including location (_to be at_, _to remain at_, etc.) and emotion (_to hate_, _to love_, etc.).

In addition to simple stative relationships, verbs can be used to convey changes of state. Following Chafe (1970) I will refer to a change of state as a _process_. For example, the sentence

Ida receives $10.00.

tells us
(1) that Ida now has $10.00
(2) that someone else had the $10.00 before;
(3) that a change has taken place from this previous state of possession to the present state.

More commonly, verbs express not simple changes of state but causal changes of state. We seem to be very interested in processes that are volitionally caused by humans and other sentient beings. Figure 2 shows the representation of the sentence:

Ida gives Sam a rose.

An agent may cause a change of state that relates to another object. Or the same person may act on both agent and experiencer of the change of state. The locational verb _move_ can be used in either way, as in the following examples:

a. Ida moved the car.
b. Ida moved to the front seat.

In both these cases the action taken by Ida is unspecified. We often don't care exactly what someone did to cause some process to occur. However, there are also verbs in which the causal action is partially or wholly specified: e.g., _walk_, _saunter_, _meander_, _stride_, _run_, _sprint_, _race_, _trot_, _jog_. (See Miller (1972) and Miller & Johnson-Laird (1976) for a more extensive discussion of the verbs of location.)

Thus, this system allows for the representation of verbs as states, changes of state, causal changes of state, simple actions, and complex cases in which specific actions cause changes of state. Further discussion of the LNR system of verb semantics can be found in the articles by Abrahamson, Gentner, Munro, Rumelhart & Levin, and Rumelhart & Norman in the Norman & Rumelhart (1975) volume.

There are certainly gaps in the system, and aspects of verb meaning that are not expressible in this simple vocabulary. Some unresolved issues are discussed later in the paper. However, the system seems plausible at the first level, and allows a fair range of verb meanings to be captured at least roughly.

At this point in the research it seemed appropriate to begin testing the psychological rightness of the system as so far stated before going on to refine it.

Psychological Tests of the Model

One advantage of psychological experimentation (or of computer implementation) is that it forces one to make explicit the assumptions underlying representation and process. At least some of the choices made can then be tested as hypotheses. Some important assumptions are

(1) a verb's representation captures the set of immediate inferences that people normally make when they hear or read a sentence containing the verb;

(2) in general, one verb leads to many inferences

(3) these networks of meaning components are accessed during comprehension, by an immediate and largely automatic process

(4) the set of components associated with a given word is reasonably stable across tasks and contexts

(5) surface memory for exact words fades quite rapidly, so that after a short time, only the representational network remains.

In testing these representations, I took a very literal interpretation of the notion of representation -- namely that the nodes and arrows in a representation correspond to the concepts and relationships that are stored when a person comprehends a sentence containing a verb. The more ferociously literal the interpretation, the better the chances of discovering counter-evidence.

Semantic overlap. One psychological criterion is that the representations should agree with people's intuitive notions of synonymity and similarity in meaning. One straightforward measure of this overlap is the degree to which people rate verbs as similar in meaning. In a study of about 60 selected verbs, I found that people's average rating of the semantic similarity between two verbs agreed very closely with the degree of semantic overlap between their representations.

A more subtle measure of psychological similarity is the degree to which people unconsciously confuse things in memory. People in a sentence-memory experiment probably try to keep their sentence traces clear. But, suppose that within a short time after hearing a verb in a sentence, a person has only the representational network of concepts and relationships, and not the surface verb. Assume further that some pieces of the memory representation may be lost or unaccessible at any time (the "fallibility of human memory" assumption). Then the more two verb representations overlap, the more likely it is that sentences containing the two verbs will be confused in memory, despite people's attempts to keep them straight. In an experiment in sentence memory, using verbs of varying semantic overlap, I found that subjects did indeed confuse the verbs in exactly the way predicted by the theory (Gentner, 1974). The correlation between the number of confusions subjects made between two verbs and the semantic overlap between the verbs, as predicated from the representations, was quite high. In fact, the correlation between representational overlap and number of confusions was slightly higher (though not significantly so) than the correlation between the number of confusions and the rated similarity between the verbs. (The similarity ratings were taken from the first-mentioned study, with a different set of subjects).

Semantic complexity. Semantic complexity refers to the number of underlying subpredicates and interconnections that make up the basic meaning of a verb. More complex meanings correspond to more specific actions or events. For example, stride is more specific than go. Its meaning contains more subpredicates. We know more having heard sentence (a) than sentence (b).

(a) Ida strode across the field.

(b) Ida went across the field.

Various researchers have looked for evidence that semantic complexity may affect comprehensibility, generally on the assumption that more complex semantic structures are harder to process (Kintsch 1974; Thorndyke, 1977). However, the results have been negative. There is no evidence that more complex words lead either to longer reaction-times or to greater processing loads than do simpler words. I believe that it's incorrect to assume across the board that complexity is psychologically hard. Some research of mine suggests that the effects of semantic complexity in memory are more particular.

Semantic Complexity and Connectivity. Although the view that semantic complexity leads to difficulty has not been supported, there is another side to the complexity issue. The additional semantic components in a complex verb may set up additional connections among the nouns in the sentence. In this case, more complex verbs should lead to a richer and more highly interwoven sentence representation, and thus to better memory for the nouns in the sentence.

Notice that this prediction derives from a fanatically literal interpretation of the verb representations: more paths in the representation means more conceptual paths in memory. This prediction is quite specific. It is not simply a question of certain complex versus simple verbs having some overall effect, but rather of complex verbs providing extra connections between the particular nouns in question. This is clearly true for Ida and her tenants in the case of sell versus give, as can be seen in Fig 3a and 3b.

I tested for this kind of improvement in connectivity in a series of experiments in sentence memory (Gentner, 1977). I read people sentences that differed in the semantic connectivity of their verbs, such as the following pair of sentences:

Ida gave her tenants a clock. (simple)

Ida sold her tenants a clock. (complex connective)

Then I gave the people the names of the characters and asked them to recall the sentences. As predicted, they were better able to recall the noun tenants when the complex connective verb sell was used then when the simple verb give was used. More semantic connections between the two nouns led to stronger memory connections.

To see the specificity of the prediction, consider a complex verb that merely amplifies the simple verb and does not add connections between the key nouns. For example, the verb mail (Fig 3c) adds the information that the method of transfer was by mailing or some such long-distance transfer. Using mail leads to more inferences (a more specific event description) than using give. However, the knowledge that the object was mailed leads to few, if any, additional connections between the agent, Ida, and the recipient, tenants. Therefore, the prediction was that use of such non-connecting specific verbs would lead to no improvement over use of general verbs in memory between the nouns.

The results were exactly as predicted: The object nouns of complex connective verbs were recalled better than those of general verbs and non-connecting complex verbs. These differences were not traceable to differences in imagery or word-frequency. Thus connectivity is beneficial to sentence memory in a very specific way.

Acquisition. There may be a more direct relationship between complexity and difficulty in children than in adults. Young children often fail to comprehend the full meanings of semantically complex terms (e.g., Bowerman, 1975; Clark, 1973; Gentner, 1975, in press). Working with the verbs of possession, I have observed that children act out the simple verbs give and take correctly before they act out the more complex verbs buy and trade. Still later they learn the yet more complex verbs buy, sell and spend. The order in which the verbs are learned is exactly the order of increasing semantic complexity. This complexity ordering can be made quite precise, since the verbs are closely related in meaning. The representation of a verb at the nth level of simplicity is properly nested within the representation of a verb at the (n+1)th level. Further, when children around 4-6 years are asked to act out sell (as in "Make Ernie sell Bert a boat.") they act out give instead (A boat is transferred from Ernie to Bert). Similarly, buy is acted out as take. They systematically act out complex verbs like simple verbs; and more surprisingly, they choose the appropriate simple verb. My interpretation, consistent with Clark's (1973) semantic features analysis, is that children learn these complex verb meanings gradually, by adding components to their partially correct representations. At any given time, the child comprehends language in terms of the components that he has so far acquired.

Semantic Integration. Another important psychological requirement is combinability. The basic notions of state, change of state, cause, and so on must be combinable into networks larger than the individual sentence. When two verbs share parts of their underlying structure, this redundancy should be utilized to combine the two representations into one discourse structure. How can we test whether this happens? One way is to arrange things so that collapsing the redundancies between two verbs should create the representation of a third verb. Then the prediction is that people should use this third verb in recall.

In a study of semantic integration, I read people short passages and tested their memory by having them fill in blanks (Gentner, 1978). Every passage contained a general verb, such as give. Half the passages also contained additional semantic information, such as the fact that the giver actually owed the money he was giving. According to the representational model, the integration of the representation of give with that of owing should have created the structure of pay. If what people have in their minds after hearing the verbs is the network representations, and if these representations are integrated during discourse comprehension, then people who heard give and owe should end up with the representation of pay. As predicted, subjects hearing the extra material falsely recalled the verb which best fit the composite structure (e.g. pay) rather than the verb actually presented.

Further Issues

I have made the assumption that a verb carries with it a set of inferences that are normally made during comprehension, as well as several supporting assumptions. This view has been fairly well supported by the research presented here, but nevertheless it seems to me an oversimplification. There remain a great many questions, some large and some small.

(1) Where should the line be drawn around a word's meaning? As Clark and Clark (1977) have put it, is word meaning more like a dictionary or an encyclopedia? The extreme of the dictionary approach would be to take a minimal contrast approach, storing with a word only enough to distinguish it from all other words. The extreme of the encyclopedia approach would be to access the entire long-term memory whenever any word is used. The question is, how to define a reasonable middle ground.

(2) What is the process of expansion into a semantic representation during comprehension?

a) Are there invariable inferences? When an incoming word is processed, is there a set of inferences (such as the set I have called the "almost-inevitable inferences" that is always made during comprehension, or is there variation in which inferences get made?

b) If there is variation, is it quantitative or qualitative? Do context and the person's interests and attention determine which inferences get made, so that there are qualitative differences in what inferences get made? Or is the difference merely quantitative, with the radius of expansion varying with the amount of attention (or energy, or interest) that the person brings to bear?

The notion of at least quantitative variation a seems hard to avoid. It is a fairly strong intuition that we process word meanings with varying degrees of energy. Further, the phenomenon of instantiation (Anderson, R.C., Stevens, K.C., Shifrin, Z., & Osborn, J.; 1977) makes it clear that a model of sentence comprehension must allow for qualitative differences in the final set of inferences stored. For example, compare the sentences

Rover ate his dinner.

Mr. Pritchard ate his dinner.

The verb eat conveys vastly different action sequences when used with different agents, though its causal change-of-state structure remains more-or-less constant. It is possible that this qualitative variation can be accounted for by simple underlying quantitative processes spreading activation. We may have to settle for a more complex model, in which some parts of a verb's meaning are almost always accessed while other inferences develop out of the interaction of the verb with its context, including its pragmatic context. In Hewitt's (1976) terms, there may be both if-added inferences and if-needed inferences. Where in this model (and whether) we want to draw a line between meaning and knowledge-of-the-world is not at all clear to me. (3) Carrying the notion of variable verb meaning still further, how does metaphorical extension work? Most common verbs can be used in several related ways. For example, consider the range of meanings that give can convey depending on the nouns it is used with:

Ida gave Sam
- a rose.
- a job.
- an heir.
- an excuse.
- a talking to.
- all his best ideas.
- the time of his life.

Clearly the subpredicate structure varies across these sentences, so much so that some might want to describe this as a collection of entirely different senses of the same word. This misses the structural similarities. Some kind of metaphorical extension of meaning seems a necessary part of a theory of verb meaning, since it is generally the verb that does most of the adjusting. A series of studies by Albert Stevens and me suggests that people faced with an odd sentence assume that some of the subpredicates normally conveyed by the verb are not meant to apply in the sentence at hand. A current project is to model the rules for which subpredicates apply in different contexts.

(4) I have so far treated nouns as nodes in the semantic representation. Clearly in order to analyze sentence interactions it is necessary to have a representation of noun meaning. Some progress been made with abstract nouns, such as kinship terms. But the truly nounlike nouns ---basic-level nouns--- resist analysis. I believe that these differences in amendability to analysis reflect differences in the kind of meaning that verbs and nouns have, and that a useful representation of concrete noun meaning may be quite different from that used for verbs, prepositions and even abstract nouns.

(5) There are several aspects of the representational scheme that need further thought. To single out one issue, consider the notion of change of state. The LNR representation represents a verb like get as conveying a change from an initial state of possession to a final state of possession. Schank's Conceptual Dependency theory would represent the entire sequence as a primitive act. Many generative semanticists have represented only the inchoative part of the chain (the change to the final state) as belonging to the assertion of the verb, considering the initial state to be more in the nature of a presupposition (e.g. Fillmore, 1966). All these positions seem to me to have merit. The LNR use of change from initial to final state allows a change-of-state verb to hook automatically with relevant state information. The use of acts as primitives captures the psychological wholeness of change. The use of the inchoative captures the intuition that people seem more interested in the results of an event --i.e. in the final state-- than in the setting state. The explicit change-of-state formats (LNR format and inchoative format) have a natural way of capturing some kinds of metaphorical extension: by substituting a different stative while preserving the rest of the verb's structure.

Summary

This work is just beginning. Neither the representations nor the processes that are assumed to operate on them come very close to capturing the subtlety of human language use. Still, the results of the experimental investigation are promising some kind of decompositional model along these lines.

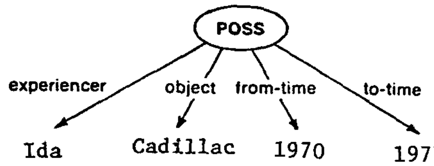

Figure 1. Ida owned a Cadillac from 1970-1977.

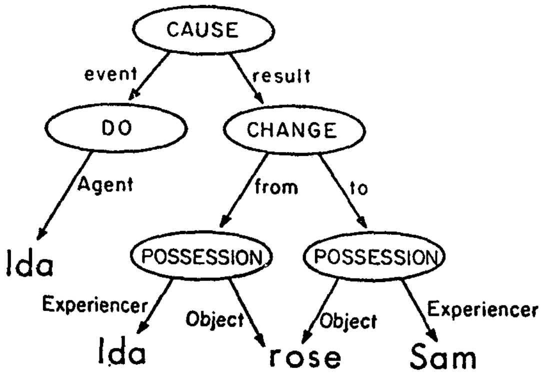

Figure 2. Ida gives Sam a rose.

Ida gave her tenants a clock

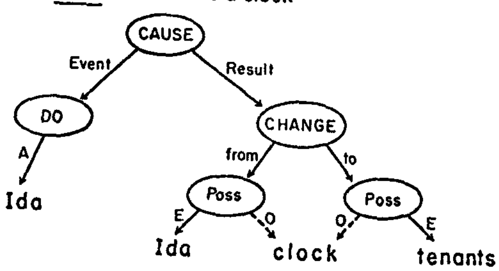

GENERAL VERB (FEW CONNECTING PATHS)
Figure 3a.

Ida sold her tenants a clock

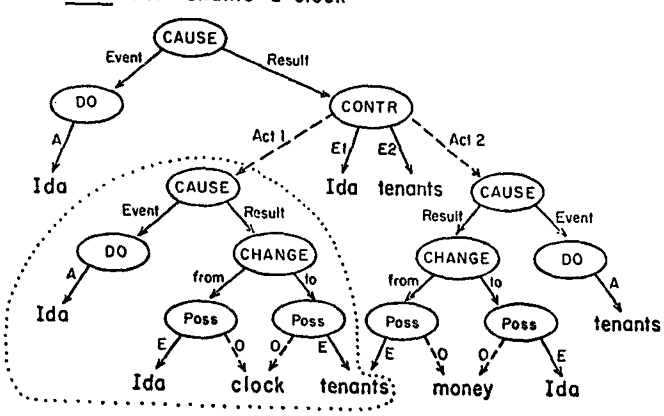

SPECIFIC VERB (MANY CONNECTING PATHS)
Figure 3b.

Ida mailed her tenants a clock

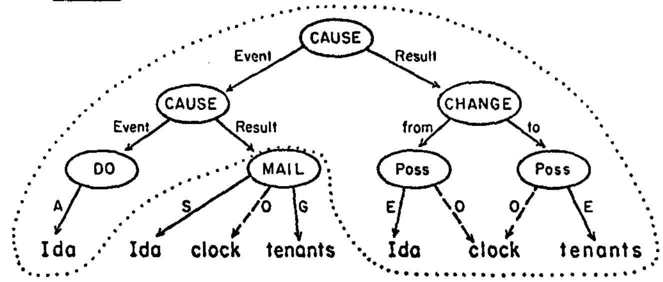

SPECIFIC VERB (FEW CONNECTING PATHS)
Figure 3c.

Footnote

1. The representational format shown here was developed by a group of researchers at the University of California at San Diego: Adele A. Abrahamson, Dedre Gentner, James A. Levin, Stephen E. Palmer, and David E. Rumelhart. The system is explained in detail in Norman & Rumelhart, 1975.

References

Abrahamson, A.A. Experimental analysis of the semantics of movement. In D.A. Norman, & D.E. Rumelhart, (Eds.), Explorations in Cognition. San Francisco: W.H. Freeman & Co., 1975.

Anderson, R.C., Stevens, K.C., Shifrin, Z., & Osborn, J. Instantation of Word Meanings in Children, May 1977.

Bendix, E.H. Componential analysis of general vocabulary: The semantic structure of a set of verbs in English, Hindi, and Japanese. The Hague; Mouton, 1966.

Bowerman, M. The acquisition of word meaning: An investigation of some current conflicts. Paper presented at the Third International Child Language Symposium, London, September 1975.

Chafe, W.L. Meaning and the structure of language. Chicago: University of Chicago Press, 1970.

Clark, E.V. What's in a word: On the child's acquisition of semantics in his first language. In T.E. Moore (Ed.), Cognitive development and the acquisition of language, New York: Academic Press, 1973.

Clark, H.H. & Clark, E.V. Psychology and Language. New York: Harcourt Brace Jovanovich, Inc., 1977.

Fillmore, C.J. Review of Bendix's Componential analysis of general vocabulary: The semantic structure of a set of verbs in English, Hindi, and Japanese. International Journal of American Linguistics, 1966, 32, Part II, No. 2. Publication 41.

Gentner, D. Towards a psychological theory of the meaning of the possession verbs. Unpublished doctoral dissertation, University of California, San Diego, 1974.

Gentner, D. Evidence for the psychological reality of semantic components: The verbs of possession. In D.A. Norman and D.E. Rumelhart, Explorations in Cognition, San Francisco: W.H. Freeman & Co., 1975.

Gentner, D. On relational meaning: The acquisition of verb meaning. Child Development, in press.

Gentner, D. Semantic integration of word meanings. Bolt Beranek and Newman Inc. Report No. 3826, May 1978. Also to appear as a Center for the Study of Reading Technical Report.

Hewitt, C. Viewing control structures as patterns of passing messages. M.I.T. AI Working Paper 92, 1976.

Lakoff, G. Adverbs and modal operators. Indiana University Linguistics Club Reprint. Bloomington: Indiana University Linguistics Club, 1970(a).

Lakoff, G. Irregularity and syntax. New York: Holt, Rinehart and Winston, 1970(c).

McCawley, J.D. The role of semantics in a grammar. In E. Bach and R.T. Harms (Eds.), Universals in linguistic theory. New York: Holt, Rinehart and Winston, 1968(b).

Norman, D. A., Rumelhart, D.E. & the LNR Research Group. Explorations in cognition. San Francisco: W.H. Freeman & Company, 1975.

Rumelhart, D.R. & Levin, J.A. A language comprehension system. In D.A. Norman & D.E.

Rumelhart, Explorations in cognition, San Francisco: W.H. Freeman & Co., 1975.

Schank, R. Conceptual Dependency: A Theory of Natural Language Understanding, Cognitive Psychology, 1972, 3, 552-631.

Schank, R.C. The structure of episodes in memory. In D. Bobrow & A. Collins (Eds.), Representation and understanding. New York: Academic Press, 1975.

Talmy, L. Semantic structures in English and Atsugewi. Unpublished doctoral dissertation, University of California, Berkeley, 1972.

Thorndyke, P.W. Cognitive structures in comprehension and memory of narrative discourse. Cognitive Psychology, 1977, 9, 77-110.