Psychological Bulletin
1984, Vol. 96, No. 2, 292–315

Only Connections: A Critique of Semantic Networks

P. N. Johnson-Laird
Medical Research Council
Applied Psychology Unit
Cambridge, England

D. J. Herrmann
Hamilton College

R. Chaffin
Trenton State College

This article examines theories that assume that semantic networks account for the mental representation of meaning. It assesses their similarities and divergencies, and argues that as a class of theories they remain too powerful to be refuted by empirical evidence. The theories are also confronted by a number of problematical semantic phenomena that arise because networks deal with the connections between concepts rather than with their connections to the world. The solution to these problems could be embodied in a new network system, but such a system would differ in both structure and function from current network theories.

In the past 15 years, a number of memory- and linguistic-processing theories have been advanced with a common approach to semantics: They have assumed that meaning is represented by a network of labelled associations. These so-called semantic network theories have stimulated considerable research, both experimental studies of categorization and computer implementations of networks. However, no systematic investigation has been made of semantic networks as a class of psychological theories. The aim of this article is to provide such an examination in order to assess the strengths and weaknesses of this approach to the psychology of meaning.

Semantic networks are used to model performance in a variety of tasks. A review could examine them from several angles (e.g., how they explain verbal learning, or how they account for the comprehension of prose). We choose to concentrate on linguistic meaning because network theories have made an important contribution to elucidating the semantic representation of words and sentences. We are concerned primarily with the empirical content of the theories rather than with the formalism of networks, although inevitably we cannot consider one without the other.

In this article, we first outline the main goals to be achieved by any cognitive theory of meaning. Second, we review the main network theories that are intended as psychological theories. Third, we consider what it is that such theories have in common and what potential constraints may be placed on the class of network theories as a whole. Fourth, we outline a variety of semantic phenomena that network theories have difficulty in accommodating. Fifth, we consider what lies at the root of these phenomena and advance a potential explanation of them. Finally, we draw some brief conclusions about the resulting status of semantic networks in relation to the goals outlined in the first section.

Goals of a Psychological Theory of Meaning

Natural language enables human beings to communicate ideas. For example, if a speaker asserts in a conversation about Christopher Columbus, "The captain of the Santa Maria thought the earth was round," listeners are able to recover the significance of the remark. This simple example of communication still defies analysis and provides us with a useful test case. It illustrates most of the major goals

The authors are grateful to Doug Hintzman, Richard Young, and two anonymous reviewers for some stimulating comments on an earlier draft of this article. The second author thanks the MRC Applied Psychology Unit and its director, Alan Baddeley, for providing a congenial atmosphere for research while he was a visitor there in 1982–1983.

Requests for reprints should be sent to P. N. Johnson-Laird, MRC Applied Psychology Unit, 15 Chaucer Road, Cambridge CB2 2EF, England.

for a psychologically plausible theory of meaning (see Bierwisch, 1970; Lyons, 1977; Miller & Johnson-Laird, 1976, p. 706).

In our view, there are four principal goals. The first is to specify the form of the mental representation of meaning. Such a representation is in the mind of the speaker and mapped into words by the speech production system. The listener decodes these words into a mental representation which, if all goes well, captures what the speaker had in mind. It is important to note the distinction, drawn by Frege (1892), between *sense* and *reference.* The reference of an expression is what it stands for in the world; the sense of an expression is that part of its meaning that concerns the way the expression connects with its reference. Hence, the reference of the expression "the captain of the Santa Maria" is Christopher Columbus, but other expressions with different senses, such as "the discoverer of America" or "the first person to sail the Atlantic," can have the same reference. Frege's informal notions of sense and reference have been replaced in formal semantics by the concepts of *intension* and *extension,* and we adopt this terminology. Thus, the intension (or sense) of a word such as *captain* determines the set of all its possible extensions (or referents), and its extension in a particular proposition, such as the one expressed by our example, is a particular individual. We are not concerned with the logical analysis of intensions and extensions, but it is crucial to bear in mind two distinct sorts of relations: the relation between the sense of one word (or expression) and another, which we call an *intensional relation,* and the relation between a word (or expression) and its referent, which we call an *extensional relation.*

The second goal of semantic theory is to explain intensional phenomena. It should account for the intensional relations of words and expressions, such as synonymy, antonymy, and inclusion (e.g., *captain* has "human" as a superordinate; the example sentence implies: "The captain of the Santa Maria did not think that the earth was flat"). It should account for the semantic properties of expressions such as ambiguity, anomaly, analyticity (i.e., truth in virtue of meaning), and self-contradiction (i.e., falsity in virtue of meaning).

The third goal is to explain extensional phenomena. The theory should account for the extensional relations between words and the world as human beings conceive it. It should elucidate the way in which speakers and listeners relate expressions to their extensions. It should also establish the way in which semantic representations capture the truth conditions of assertions and the senses of questions and commands. A major puzzle is how a semantic representation captures the fact that assertions are true with respect to an infinite number of different situations. Even the simple remark about Columbus is true of many different possibilities: He might have thought that the world was an ellipsoid, cylindrical, pear shaped, and so on. A picture may be worth a thousand words, but a proposition is worth an infinity of pictures.

The fourth goal is to explain the inferences that people draw in virtue of the meanings of words. That is, the form of semantic representations should dovetail with whatever machinery is used to make verbal inferences.

There are undoubtedly other goals for semantic analyses, but these four are perhaps the most crucial for a psychological theory. They all concern what is computed in various semantic tasks, and obviously the theory must also account for how the tasks are carried out (i.e., for the various mental processes underlying them). The theory should specify an effective procedure (or algorithm) by which the mind constructs semantic representations during comprehension and an effective procedure by which semantic representations are mapped into words during speaking. These procedures presumably depend on a grammar, a lexicon, knowledge of the world, and the ability to make inferences. The theory should also specify procedures for evaluating intensional properties and relations, because speakers are able to judge whether an expression is ambiguous, anomalous, or meaningful (e.g., Steinberg, 1970). They are also able to make judgments of synonymy and antonymy (e.g., Herrmann, Chaffin, Conti, Peters, & Robbins, 1979). Finally, the theory should specify procedures for representing and evaluating extensions and truth values, because speakers can determine what an expression refers to and ascertain whether assertions are true (Clark & Clark, 1977).

In short, a psychological theory of meaning should explain how meaning is mentally represented, how expressions are intensionally related one to another, how they are related to the world as the mind conceives it, and how their semantic representations enter into inferences. In due course, we examine how well semantic networks meet these explanatory goals, but first we outline this class of theories.

Semantic Network Theories

Early Semantic Networks

Semantic networks are associative theories framed for computers. Associationists from Aristotle to the present day have assumed that there can be an association from one word to another. The set of associations to a given word will contain many members that will differ in strength, and many of the words in the set will themselves be associatively related. It is therefore natural to think of words as associated by a network of undifferentiated links varying in strength (see, e.g., Deese, 1962, 1965; Kiss, 1967; Walpole, 1941; Warren, 1921). Although such a network may be explanatorily useful for studies of word association, it is obviously a poor instrument for semantics because a mere associative link from one word to another tells one nothing about the intensional relation between the words. For example, *black* is strongly associated with both *white* and *night*, but the intensional relation between *black* and *white* is very different from the intensional relation between *black* and *night*.

The decisive step in converting an associative network into a potential semantic machine was the introduction of labels on the links between words. The importance of representing the relation between entities was recognized by Selz in 1913 (cited by Humphrey, 1951), but Selz lacked an appropriate language in which to make his theory wholly explicit. With the advent of programming languages, and particularly list-processing languages such as IPL and LISP, the way was clear for the invention of semantic networks. An expression in LISP is either a *list* or an *atom*, and an atom is either a number or a symbol, such as *poodle*. LISP allows one to set up a whole series of property names and their specific values associated with a symbolic atom such as *poodle:*

Superordinate: dog

Size: small

Hair: curly

In describing such information, it is natural to represent it in a graphical form:

$$\begin{tikzcd} \text{Dog} & \text{Poodle} \arrow[l, "\text{Superordinate}"'] \arrow[r, "\text{Size}"] \arrow[d, "\text{Hair}"'] & \text{Small} \\ & \text{Curly} & \end{tikzcd}$$

Early semantic networks (Quillian, 1968; Raphael, 1968) were probably inspired in part by this feature of IPL and LISP.

Quillian's (1968) theory anticipated most of the features of subsequent semantic networks. He assumed that memory for meaning is no different from memory for perceptual or other nonlinguistic information, and he postulated a semantic network as a model for lexical memory. The network is composed of links between two sorts of nodes: *type* nodes, which represent concepts, and *token* nodes, which represent instances of concepts by virtue of the links to their respective type nodes. The meaning of a word is defined by an initial configuration of token nodes attached to the type node representing the word, and each of the token nodes is linked to its respective type node. Figure 1 presents a simplified example of such a network, which represents one sense of the verb "to plant." As the figure illustrates, the theory proposes five sorts of links: (a) From a token to its type node (as shown by the dotted lines leading from tokens out to their respective types, which are omitted from the figure); (b) from a type to a superordinate token (e.g., the link from *plant* to *put*); (c) from a token to another token representing a modifying property (e.g., the link from *put* to *for* to represent the fact that seeds are put in the ground for them to grow); (d) from a token to a conjunction or disjunction of tokens (e.g., from *object* to *seed, plant,* or *thing*); (e) from a token to two other tokens so as to act as a label on a link between them (e.g., *object* and *earth* are related by *in* because the object

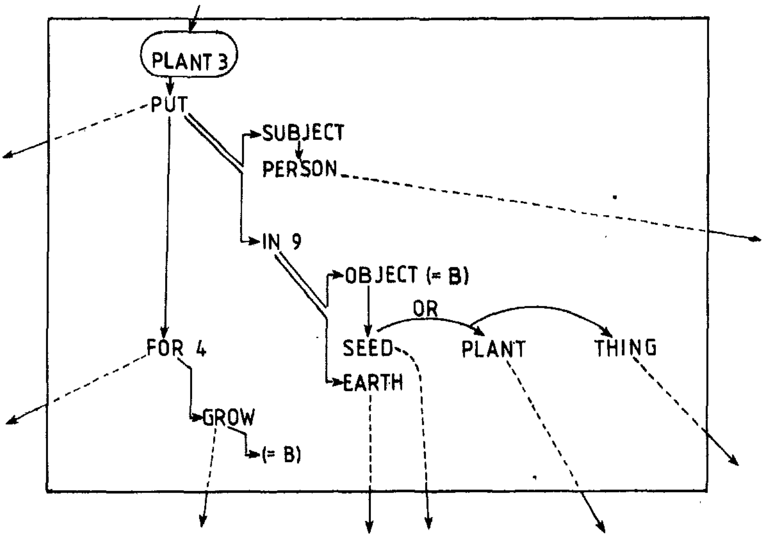

Figure 1. A semantic network representation of one sense of the verb "to plant." (From "Semantic Memory" by M. R. Quillian, in *Semantic Information Processing* [p. 236] edited by M. L. Minsky, 1968, Cambridge, MA: MIT Press. Copyright 1968 by MIT Press. Reprinted by permission.)

is put in the earth). The amount of information in a network is potentially so vast that Quillian assumed that facts are stored explicitly only if they cannot be generated from the network. Hence, general information need be represented only at a superordinate level without being attached to all the subordinate nodes to which it applies. For example, a poodle is a dog and a poodle is an animal, but because all dogs are animals, it is parsimonious to use superordinate links in the network:

Poodle → Dog → Animal

There is no need to represent the fact that a poodle is an animal, because this information can be inferred from the network by tracing through the pathway from *poodle* to *animal.*

Quillian introduced various "tags" that could be attached to nodes. They include a number representing the degree to which a property represented by one node applies to an entity represented by another node, a number representing the "criteriality" of a token in defining a type, and special symbols that indicate the role of other words in defining a concept. In this last case, Quillian used symbols to show that one word could be the subject or the direct object of the current word, or that a word could be modified by the current word. Thus, as Figure 1 shows, one sense of *plant* calls for a person, as subject, putting seeds, plants, or other things, as object, into the earth. Quillian suggested that this machinery could be used to select the appropriate meaning of the verb, which is otherwise ambiguous, from the linguistic context in which it occurs. He also suggested that the labels of Fillmore's (1968) Case grammar might be more appropriate than these standard grammatical terms.

Such a theory about representation becomes a full-fledged theory of performance only when processes for using the representation—for setting it up, interrogating it, and drawing inferences from it—are specified. Indeed, as Quillian suggested, certain aspects of meaning may be represented by such processes rather than directly in the network. A natural assumption, however, is that when people evaluate the relation between two concepts, they do so by searching through the network for a path between them. Quillian wrote a computer program that operates in this way. The program simulates parallel processing by making a breadth-first search along the links radiating out from the two type nodes activated by the input words. At each node, the program leaves tags specifying both the immediately preceding node in the search and the original starting node. Thus, the search moves outward in a constantly expanding spread of activation. An intersection occurs when the search from one starting point encounters a node with the tag from the other starting point. The program then evaluates the path between the starting nodes in order to establish the nature of the intensional relation between them.

Quillian's model aroused interest among psychologists because it could be used to predict how long it takes to verify different sorts of statement. Collins and Quillian (1969, 1972) assumed that the greater the length of the path that has to be traversed to establish an intersection, the longer the task should take, and they reported some experimental evidence corroborating this assumption. They found that subjects take longer, for example, to evaluate the assertion that "a poodle is an animal" than to evaluate the assertion that "a poodle is a dog." However, later research has revealed a more complicated picture. In particular, some results run counter to the proposal that properties are always stored at the most economical superordinate node (Conrad, 1972). Other results suggest that the relative sizes of sets may be more critical than path length (Landauer & Freedman, 1968). Moreover, some instances of a concept are more prototypical than others, and assertions about such typical instances (e.g., "A robin is a bird") are verified faster than assertions about atypical instances (e.g., "A penguin is a bird"; Rips, Shoben, & Smith, 1973; Rosch, 1973; Smith, Shoben, & Rips, 1974). In short, cognitive economy is not invariable, and path length alone does not always determine verification time.

The semantic network as we have described it so far constitutes a theory about the organization and processing of the mental lexicon: Representations of words are stored in a network, and the semantic relations between words are represented by labelled links between the items in the network. Semantically related words have shorter paths between them, and it is supposed to take less time to traverse these paths than it does to traverse those between more remotely related words. Analytic assertions, which are true in virtue of the meanings of words (e.g., "A poodle is an animal"), and self-contradictory assertions, which are false in virtue of the meanings of words (e.g., "A poodle is a car"), are evaluated by checking the relation that is represented in the network.

The next major step in the evolution of network theory was to show how the meaning of any sentence could be represented as a semantic network. A number of different theories were proposed to do this job, but they have in common the following basic idea of Quillian's. If a sentence uses a verb to establish a relation between the entities denoted by noun phrases, then the semantic representation of the sentence can take the form of a small-scale semantic network that captures the relation (expressed by the verb) between the entities (denoted by the noun phrases). The full semantic representations of the words in the sentence are, of course, encoded separately in the main network corresponding to the mental lexicon. The interpretative process setting up the representation of the meaning of a sentence must contain a mechanism, such as the one previously described, for selecting the appropriate sense of an ambiguous word as a function of the context in which it occurs.

Network theories differ principally in their configurations of links and in the types of label that they use. These differences to some extent echo those in the linguistic theories current when the networks were formulated. We now examine the major psychological network theories in more detail, concentrating on the differences between them and on their contrasts with Quillian's seminal formulation.

Human Associative Memory (HAM) Network

Human Associative Memory (HAM) is a computer model of long-term memory developed by J. R. Anderson and Bower (1973). Like Quillian, these authors assumed that both linguistic and perceptual information is stored in the form of abstract propositional representations (i.e., structures in a network). This assumption was motivated by two phenomena: (a) that people primarily remember the gist of sentences, not their details verbatim; and (b) that subjects can make inferences—it was presumed that they were based on rules of inference that operate on propo-

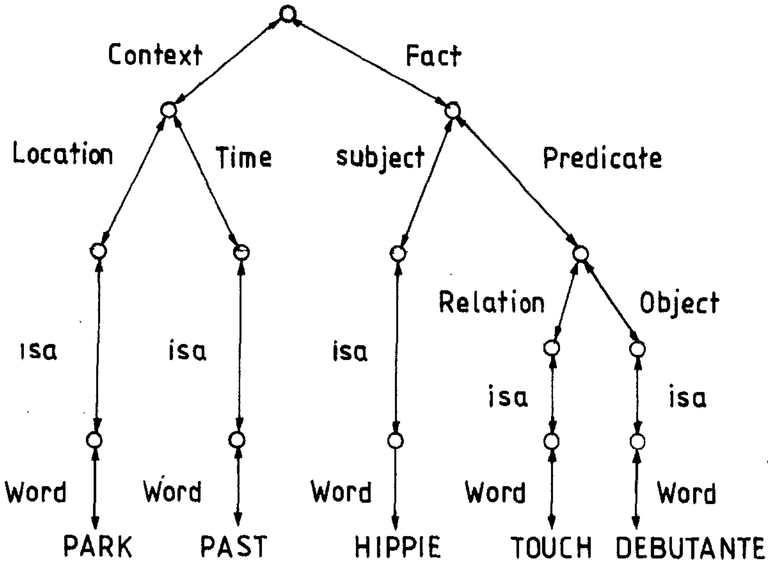

Figure 2. The HAM representation of "In a park a hippie touched a debutante." (ISA denotes set membership. From *Human Associative Memory* [p. 67] by J. R. Anderson and G. H. Bower, 1973, Washington, DC: Winston. Copyright 1973 by Winston. Reprinted by permission.)

sitional representations. J. R. Anderson and Bower devised a simple parser that maps a limited subset of English into network representations. Figure 2 presents their network for the sentence, "In a park a hippie touched a debutante." Each word has a corresponding concept represented by a circle. All the semantic information is represented by the configuration of the links and the labels on them; the nodes themselves have no semantic labels. The sentence is represented in two parts: the *context* ("in a park") and the *fact* ("a hippie touched a debutante"). The ISA label denotes set membership, and the remaining labels are reminiscent of those to be found in the standard theory of transformational grammar: subject, predicate, object. As in Quillian's theory, the concept nodes are taken to exist already in memory, and the representation of the sentence accordingly depends on them.

In addition to the features illustrated in Figure 2, HAM also contains machinery for coping with some aspects of quantifiers (e.g., it has symbols denoting universal quantification and implication). The theory also assumes that a process of pattern matching searches memory for the existence of a given proposition. J. R. Anderson, however, concluded that HAM is incomplete, and he created a new network theory, which we describe shortly.

Lindsay, Norman, and Rumelhart (LNR) Model

Over a number of years, Lindsay, Norman, and Rumelhart developed a theory of long-term memory and comprehension (see, e.g., Rumelhart, Lindsay, & Norman, 1972; Norman, Rumelhart, and the LNR Research Group, 1975; Lindsay & Norman, 1977). Their system again decomposes into a semantic network, processes that operate on it, and a parser that maps sentences into a network representation. The semantic network comprises labelled nodes (the only items that are directly addressable in the network) and labelled links between them. There are four sorts of nodes in an LNR network: concept, primary, secondary, and event nodes. A *concept* node represents a concept, and the main links used in defining it are: ISA (for set inclusion), IS (for attributing properties), and HAS (for attributing proper parts and so on). Some concepts do not have any simple linguistic correlates, but a *primary* node does relate directly to a word in the language, and

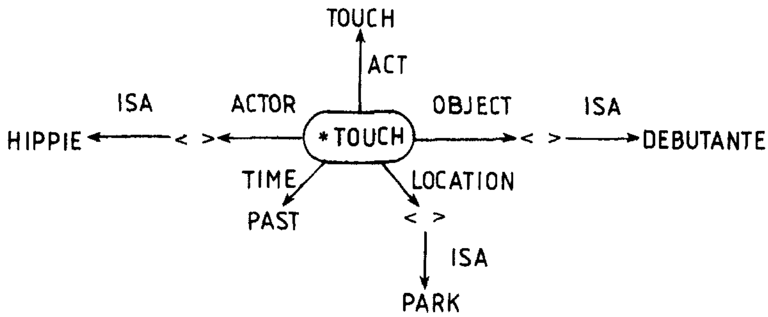

Figure 3. The LNR representation of "In a park a hippie touched a debutante." (ISA denotes set membership.)

may have attached to it a definition of the word (in the network format). A *secondary* node represents a specific use of a primary node (e.g., in representing the meaning of a specific sentence). Thus, the distinction between a primary and a secondary node is akin to Quillian's distinction between a type node and a token node. Events have a special status in the network and are represented by *event* nodes that are linked to the actions, actors, and objects involved in the event. A specific event is encoded by a secondary node with a link, labelled ACT, to the corresponding primary node. Thus, concept and event nodes are distinguishable by whether an ISA or an ACT link is attached. Figure 3 shows an LNR network representation for the sentence, "In a park a hippie touched a debutante." As the figure shows, the labels on links from nodes are based on the cases (e.g., agent, object, recipient) of Fillmore's (1968) Case grammar. In addition, there are a number of operations that increase the power of the system: operations for forming various connections between propositions (e.g., conjunction), operations for introducing quantification, and operations that generate new relations from old ones (e.g., to capture adverbs such as *slowly*).

Like Quillian, the LNR group assumes that in the lexical network general properties are stored at just one node (e.g., the fact that canaries have wings is recoverable from the fact that a canary is a bird and birds have wings). However, this generic information is treated as a set of values that can be assumed by default (e.g., typically, birds fly) rather than as necessarily true. Hence, the system can absorb exceptions (e.g., penguins do not fly) without risk of inconsistency.

In their later work, the LNR group argued that a representation such as Figure 3 is superficial, and that a deeper level of representation is required in order to capture the full meanings of sentences. These deeper representations decompose the superficial ones into their appropriate semantic primitives. Figure 4 presents such a primitive meaning structure, which is obtained by replacing items in the superficial network by their network definitions.

Glass and Holyoak's "Marker" Theory

Glass and Holyoak (1974/1975) proposed a network theory of the lexicon, which is in essence a revision of the original Quillian model designed to accommodate some experimental results. They suggested that semantic markers, similar to those postulated by Katz (1972), form a structure that underlies the intensional relations between words. Figure 5 presents an example of such an associative structure. Each noun is mentally associated with a defining marker, which is supposed to represent an abstract concept roughly equivalent to possessing the essential properties of X, where X is the noun in question. (How an uninterpreted symbol can serve this semantic function is not explained, and merely connecting it to other uninterpreted symbols would not seem to suffice.) There are other markers, such as *pet–canary,* that are not directly associated with a word. The links between markers represent relations between concepts. For example, the links

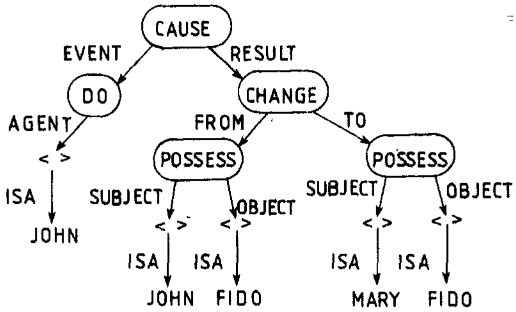

Figure 4. The LNR primitive meaning representation of "John gave Fido to Mary." (ISA denotes set membership. From "A Language Comprehension System" by D. E. Rumelhart and J. A. Levin, in *Explorations in Cognition* [p. 193] edited by D. A. Norman, D. E. Rumelhart, and the LNR Research Group, 1975, San Francisco: Freeman. Copyright 1975 by Freeman. Reprinted by permission.)

from *avian* to *feathered* and *animate* capture the fact that *avian* stands for the set: {(*animate*), (*feathered*)}. The links accordingly denote set inclusion. The chief innovation is that there are two types of intersection of links: *consistent* and *inconsistent* intersections. (The authors use the terms contradictory and noncontradictory, which we have eschewed because contradictory has the unfortunate connotation that invariably one link must be true and the other false, whereas in fact both links can be false of a given entity.) Something that is animate can be avian or canine, but it cannot be both, and so these two links are inconsistent with one another. Something that is animate can, however, be both avian and a pet, and so these two links are consistent with one another. Links that have the same Greek symbol in Figure 5 are accordingly inconsistent with one another, whereas links that have different Greek symbols are consistent with one another. Because the links represent the transitive relation of set inclusion, consistency or inconsistency can arise at a relatively remote point in the network (e.g., a canary cannot be a collie, because there is an inconsistent intersection at the animate node). This scheme is a direct reflection of Katz's (1972) treatment of antonymy. Inconsistencies can also arise in quite a different way. An assertion such as "All birds are canaries" does not yield inconsistent links but one can judge it to be false by thinking of another species of birds, such as robins, that are in an inconsistent relation with canaries. In other words, a putative generalization can be falsified by a counterexample.

Glass and Holyoak (1974/1975) allowed that a network need not be strictly hierarchical (as did Quillian and others). There may be, for example, a direct link from *canary* to *animate*. As they recognized, however, the introduction of this sort of possibility releases the network from any formal constraints on its configurations with the ensuing danger that it becomes empirically vacuous. They therefore tried to use empirical results to place constraints on the theory. In particular, they used a sentence-completion task to provide evidence about the order in which links are searched. Subjects had to complete such fragmentary sentences as "All canaries are . . ." either so as to produce a true assertion or else so as to produce a false one, and the relative frequency of the different responses

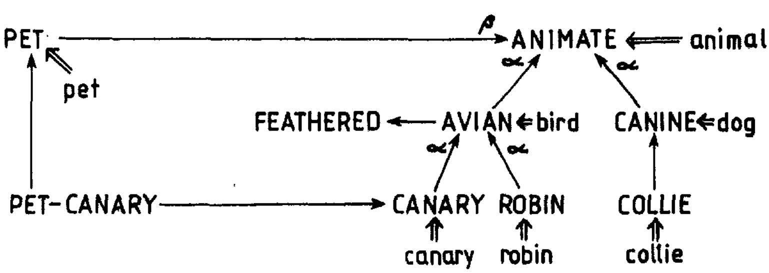

Figure 5. An example of an associative network for semantic markers. (Words are represented in lowercase letters, and markers in capital letters. Links labelled with the same Greek letter are inconsistent with one another, whereas links labeled with different Greek letters are consistent with one another. From "Alternative Conceptions of Semantic Memory" by A. L. Glass and K. J. Holyoak, 1974/1975, *Cognition, 3,* p. 319. Copyright 1974/1975 by the Associated Scientific Publishers. Reprinted by permission.)

was treated as an index of search order. This order predicts the latencies of evaluating analytic assertions, because subjects should respond faster to an assertion of the form "All *A*s are *B*s" than to one of the form "All *A*s are *C*s" if the A → B link is searched prior to the A → C link. (Nothing hinges on the notion that one link is literally searched before another: In a parallel search, one link could be traversed faster than another.) The true completions do indeed predict the latencies of evaluating true assertions. The trouble with this empirical constraint, however, is that most theories, whether or not they are based on networks, would predict the same correlation: It is merely Marbe's law as applied to semantic judgments (see Johnson-Laird, 1974). What is more striking are the data from the false sentences. Those that arise from an inconsistency of links (e.g., "All birds are dogs") were evaluated with a latency that was predicted by the frequency with which the predicate occurred in the false completions. Such pairs as bird and dog are judged as being closer in meaning than pairs that are produced relatively infrequently and evaluated relatively slowly. The results seem to run counter to the hypothesis that false assertions containing subject and predicate terms that are judged to be semantically related take longer to evaluate, but the experiment has been criticized on methodological grounds (see McCloskey & Glucksberg, 1979; Shoben, Wescourt, & Smith, 1978). False sentences that arise from the existence of counterexamples rather than from inconsistent links (e.g., "All birds are canaries") require the subject to discover a subset of the subject (such as "robins") that is inconsistent with the predicate. The latency of the evaluation should therefore depend not on the frequency with which the sentence itself is produced as a false completion but on the frequency with which the counterexample category occurs in true completions of "Some birds are. . . ." This prediction was also borne out by the results.

Glass and Holyoak's theory is deliberately restricted to relations of set inclusion, and the strong point of their model is its detailed predictions of the times it takes to reject the two kinds of false set-inclusion statements. Other semantic models, together with other interpretations of the data, are restricted to the single prediction that the more related in meaning the subject and predicate of a false statement are, the longer it will take to reject that statement (McCloskey & Glucksberg, 1979; Smith et al., 1974). The principles of the Glass and Holyoak model can be readily incorporated within the more general Quillian-based network theory, and, in part, the next theory that we consider was designed for this purpose.

Collins and Loftus Theory

Collins and Loftus (1975) pointed out a number of common misconceptions about Quillian's original theory, and proposed a number of extensions to account for various empirical phenomena. Some of these assumptions concern the way in which activation spreads through the network; others concern the organization of the network itself. Collins and Loftus drew a clear distinction between the lexical network, which is mainly organized on the basis of phonemic similarity, and the conceptual network, which is organized on the basis of semantic similarity. Semantic similarity depends on the number of properties that concepts share, and thus the number of links between them. It is distinct from semantic distance, which is the shortest distance between the relevant nodes: Two nodes may be close, yet not highly related (e.g., *cherries* and *fire engine* are relatively close because they adjoin *red,* but they have no other links in common). One of the major consequences of the new theory is that if, for example, *vehicle* is primed, then all the different types of vehicles will be primed and will prime each other, whereas if *red* is primed, then *fire engine* and *cherries* will be primed, but there will be much less mutual priming because they have no other links in common.

The authors also outlined a revised version of the decision process postulated by Collins and Quillian (1972). To decide whether two concepts match, sufficient evidence must accumulate to exceed either a positive or a negative threshold. Evidence accumulates from different pathways, with pieces of positive and negative evidence cancelling each other out in an essentially Bayesian way.

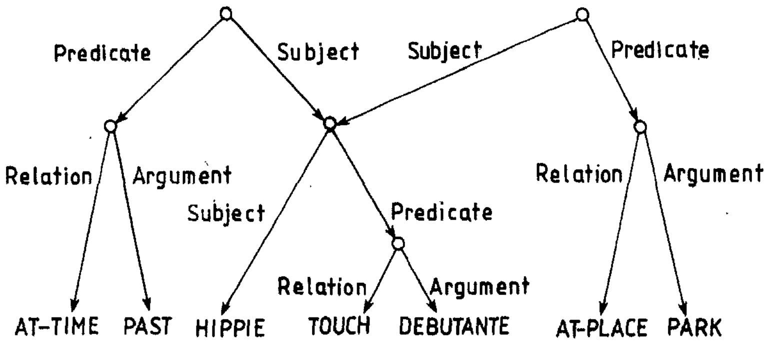

Figure 6. The ACT representation of "In a park a hippie touched a debutante." (Following J. R. Anderson, the diagram has omitted the nodes representing words. From *Language, Memory, and Thought* [p. 166] by J. R. Anderson, 1976, Hillsdale, NJ: Erlbaum. Copyright 1976 by Erlbaum. Reprinted by permission.)

Positive evidence consists of paths that establish that one concept is a superordinate of the other, or that they share a common property (depending on how criterial that property is with respect to each concept), or that one concept has a property of an instance of the other (the Wittgenstein strategy, because it occurs with concepts that depend on family resemblances among their instances rather than on a set of common elements, e.g., different sorts of games). Negative evidence consists of paths that establish that one concept is not a superordinate of the other, or that they have properties that mismatch (depending again on criteriality), or that one concept lacks the properties of instances of the other. In addition, following Holyoak and Glass (1975), negative evidence is provided by establishing that the two concepts are mutually inconsistent subordinates of the same superordinate, or by finding a counterexample to the putative relation. Although Collins and Loftus (1975) postulated only these kinds of evidence for categorization tasks, they do allow that there may be other kinds, particularly for answering complicated questions.

One consequence of these various revisions is that the Collins and Loftus model has much the same strengths and weaknesses already attributed to the Glass and Holyoak model. Because the network is no longer strictly hierarchical, there are no strong formal constraints on its configurations. In contrast the model can accommodate the experimental results that had been interpreted as contrary to network theory. In particular, the theory is now compatible with the existence of prototypes. As we have already pointed out, path length within the network does not always determine verification time. Rather, set-inclusion statements about typical instances of a category (e.g., "A robin is a bird") are verified faster than statements about atypical instances (e.g., "A penguin is a bird"; Rips et al., 1973; Rosch, 1973; Smith et al., 1974). The phenomenon is now explicable on the grounds that an atypical instance elicits negative evidence to a greater degree than does a typical instance.

ACT

J. R. Anderson (1976) developed ACT as a more advanced theory of representation than its precursor HAM. The theory assumes that only a portion of the network is active at any moment, with activation spreading from one node to another, although there is a general dampening of activation to prevent it from getting out of control. Links in the network can vary in their strength, and the spread of activation depends on the current strengths of the various links. The machinery for manipulating the network consists of a production system in which several productions can be applied at the same time. A typical representation of a sentence in the ACT system is shown in Figure 6. As the figure illustrates, sentences are no longer divided into facts and contexts; the notation is now more uniform and closely related to the standard theory of transformational grammar (Chomsky, 1965; cf. J. R. Anderson, 1976, p. 158). Anderson rejected the use of Case-grammar labels on the grounds that they lack coherent semantics. Unlike the LNR system, ACT does not decompose the meanings of words into a network of primitives, but rather analyzes their semantics by using productions. The theory therefore resembles the proposals of both Kintsch (1974) and Fodor, Fodor, and Garrett (1975) in that inferences are based on rules rather than on decomposition into a network of semantic primitives. Finally, ACT contains more powerful machinery for coping with quantification, although much of it is ad hoc (e.g., the sentence "All philosophers have read all of Plato's books" cannot be directly represented in the network format although it can be encoded as a production). The expressive power of ACT accordingly depends on both the network and the production system.

What is a Semantic Network?

Assumptions Underlying Network Theory

There is no doubt that network theories have made a significant contribution to the study of semantics. They have inspired a considerable amount of experimentation, which has revealed new empirical phenomena, and they have been modeled in many computer programs, including systems devised purely for the purposes of artificial intelligence (see, e.g., Bobrow & Collins, 1975; Findler, 1979). There is such a variety of network theories, however, that a skeptic might suspect they have nothing in common apart from the name. In fact, there appear to be four main assumptions that they share which we now try to make explicit.

First, network theories are designed primarily to elucidate intensional relations, in particular, the relations between the meanings of words. They embody no general principles concerning extensions.

Second, and a corollary of the previous point, semantic networks are constructed on the assumption that the representation and evaluation of intensional relations can be considered independently from extensional relations (e.g., dogs can be represented as a subset of animals without worrying about how to characterize the extensions of the two terms).

Third, comprehensive network theories are based on a formalism containing three components: a parser, a semantic memory consisting of a network of links between nodes, and a set of interpretative processes that operate on the network. Intensional relations between words are represented within the semantic network by labelled links between the nodes. The parser uses this information to construct network representations of sentences. The interpretative processes carry out such tasks as updating the semantic network, making inferences from sentences, and searching the semantic network to establish the intensional properties of expressions and intensional relations between them. These aspects of intensions are a direct consequence of what is encoded in the network system. In particular, the analyticity of a sentence follows directly from information in the network (e.g., the truth of an assertion such as "Canaries are birds").

Fourth, there is a general, although not absolute, commitment to parsimony. If information about the meaning of a word can be inferred by traversing links, then it is not redundantly specified in the network. Hence, general facts can be stored at the level of a superset rather than for each subset to which they apply.

These four assumptions are all that we can discern as common to network theories, but, as we now show, some of the apparent diversity is superficial.

Differences Between Network Theories

The best way to distinguish between the superficial and the more profound differences among the theories is to imagine that a theorist is setting up a new network theory. Certain questions will inevitably arise about its organization and contents. Some of these questions are impossible to answer from within the framework of network theory. Others are less problematical, and their answers will help to place some further constraints on the set of possible network theories. There are many decisions that the theorist developing a new system has to make; we

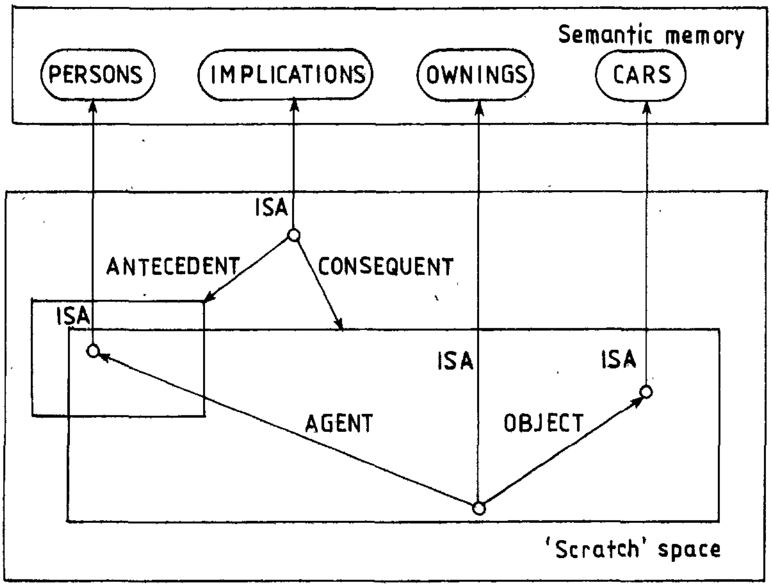

Figure 7. A "partitioned" semantic network representing the sentence, "Every person owns a car." (If a node occurs in both the antecedent and the consequent space of an implication, it represents a universally quantified variable. ISA denotes set membership. From "Encoding Knowledge in Partitioned Networks" by G. G. Hendrix, in *Associative Networks: Representation and Use of Knowledge by Computers* edited by N. V. Findler, 1979, New York: Academic Press Inc. Copyright 1979 by Academic Press Inc. Reprinted by permission.)

consider nine of the more important ones here.

1. The theorist must select an appropriate subset of a language to represent within the network. This choice will determine the expressive power needed by the network. A crucial decision is whether to admit quantified assertions, such as "No one voted for some of the candidates." Early theories did not attempt to cope with quantifiers; subsequent theories admitted them but were unable to represent them properly, because as Woods (1975) pointed out, they were directly linked to noun phrases as though they were adjectives. One problem with this practice is that it fails to specify the scope of quantifiers, and it is therefore impossible to represent the correct meaning of sentences containing more than one quantifier. Thus, the sentence "No one voted for some of the candidates" means that there are some candidates for whom no one voted. The quantifier "some" has a wider scope than the quantifier "no one" rather than vice versa. Later theories certainly admit quantification, provided that the quantifiers range over individuals, that is, the networks are equivalent in power to the first-order quantificational calculus developed in formal logic.

2. The theorist has to decide how much expressive power to put into the network and how much to put into the interpretative processes that operate on it. This decision is again crucial for quantified assertions. Some network theories place the interpretative burden of quantifiers on the processes that operate on the network (e.g., ACT). However, partitioned semantic networks of the sort devised as an exercise in artificial intelligence (Hendrix, 1979) can represent quantified assertions more directly. Figure 7 shows such a representation for the sentence, "Every person owns a car." Another method that could be implemented is to use numerical indexes on quantified noun phrases to specify their relative scopes (see Johnson-Laird, 1970). Among psychological theories, there is still no consensus about whether all sentences, including quantified assertions, should be directly representable in the network, or by processes that act on the network, or by some admixture of the two components. Thus, the

theorist devising a new system here lacks constraints to guide its development, because no principles have yet been formulated to determine what goes into the network and what goes into the processes.

Because networks have been developed without a principled division of labor between representation and process, we suggest that the following constraint might be adopted. The information in the network should represent what people remember about the meanings of sentences (i.e., it should encode the propositions that are expressed by discourse). Hence, the processes that operate on the network should play the same sort of semantic roles for all sentences, and there should be no difference from one sentence to another in terms of the type of information that is represented in the network. It follows that quantified and other complex sentences will call for the use of partitioned networks or some other structures of comparable power (cf. Brachman, 1979).

3. The theorist has to decide on the sorts of configuration that can occur within networks representing sentences. Some systems are restricted to a binary division of links from each node (e.g., HAM), and others allow an arbitrary number of links from each node (e.g., the LNR system). In principle, the semantic representation of a sentence could be made within a strictly binary network in which each node has only two links emanating from it. Certain predicates in natural language express relations between more than two arguments (e.g., "John sold Mary a dog for $10." However, links can represent functions that map two-place relations into relations with a greater number of arguments, and then such sentences can be accommodated, because a label on one of the links in the initial binary division can denote such a function (see Figures 1, 2, and 6 for variations on this theme). The same issue arises in the development of data base management systems and might be resolved in a similar way (cf. Bilger, 1979).

4. The theorist will have to decide whether to use labels on nodes as well as on links (e.g., Quillian) or to use labels only on links (e.g., ACT). This question is largely a matter of computational convenience. To represent the fact that poodles are dogs, it makes no semantic difference, given the appropriate processing system, whether one uses the orthodox convention:

$$\text{Poodle} \xrightarrow{\text{ISA}} \text{Dog}$$

or does without node labels:

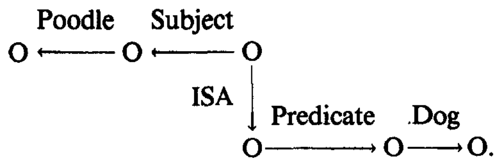

5. The previous decision is related to the question of whether all links in the network should have the same semantic force. As Woods (1975) pointed out, there is often no single interpretation for the links in the system. Such links as:

$$\text{Poodle} \xrightarrow{\text{Size}} \text{Small}$$

$$\text{Poodle} \xrightarrow{\text{ISA}} \text{Dog}$$

$$\text{Bite} \xrightarrow{\text{Agent}} \text{Poodle}$$

are semantically heterogeneous. Size denotes a function that yields a value for a given argument, ISA denotes a relation between two sets, and agent denotes part of the relation between a verb and its arguments. From within the framework of network theory, however, there is no need to impose a uniform interpretation on links. Moreover, if there were such a need, a simple strategem is available to the theorist: Each link could be treated as denoting a function, and links that normally express relations could be treated as denoting functions that return the truth values of relations, depending on whether the relation holds between the relevant entities.

6. A decision must be made about whether to use standard grammatical labels such as *subject* and *object* on the links for sentence networks (e.g., ACT), Case-grammar labels such as *agent, instrument,* and *recipient* (e.g., the LNR system), or some other idiosyncratic set of labels (e.g., HAM). This issue cannot be resolved within the framework of network theory. If a theorist's aim is to account solely for intensional relations, then it makes little

difference which system is adopted provided that it is used consistently.

7. What should be the set of semantic labels that are used in the network? Again, there is no single uniform and generally accepted set of labels other than ISA, HAS, and the type-to-token label. Intensional relations place few constraints, if any, on the particular set of labels that is selected. Nevertheless, a choice of labels may run into problems (see Chaffin & Herrmann, 1984), and unfortunately semantic network theory contains as yet no principles for constraining the set of primitives.

8. It is necessary to decide whether inferences should be made by decomposition into semantic primitives within the network (e.g., LNR) or rules of inference ("meaning postulates") that operate on the network representation (e.g., ACT). Once again, this issue cannot be decided within the framework of network theory, and some theorists claim that the choice is immaterial (J. R. Anderson, 1976). We argue later that networks fail to represent the crucial information required for inferences; it follows that neither approach is likely to be completely successful.

9. Performance in evaluating intensional properties and relations should be explicable in terms of computations carried out on the network (e.g., the latency of a response should reflect such factors as the distance that has to be traversed within the network, the order in which links are searched, and the nature of the subject's task). The theorist has to decide how these processes are to operate. However, there appear to be no constraints on the sort of processes that can be invoked to model performance in semantic tasks. Typically, processing calls for a search for a path from one node to another, but computations on the information represented in the network could take any form that the theorist desires. Indeed, most network theories seem to have the power of a Universal Turing machine (i.e., the processes they invoke are sufficient to compute anything that can be computed at all), granted Turing's thesis—yet to be falsified—that any function for which there is an effective procedure can be computed by a Turing machine (see, e.g., Rogers, 1967). J. R. Anderson (1976) has in fact shown that the ACT system is equivalent in power to a Universal Turing machine, and doubtless the equivalence could be proved for other network systems, especially those that employ augmented transition network (ATN) parsers (e.g., LNR),

These nine decisions face anyone developing a network theory. They constrain the nature of the resulting system to some extent, and we have suggested some additional factors that should be considered in making the decisions. Yet the constraints remain too few and too weak; current network systems are too powerful. Any particular network theory usually yields clear empirical predictions, but the class of network theories is not constrained: If networks can compute anything, then plainly as a class of theories they are almost empirically vacuous. The moral of Turing machine equivalence is that any experimental result that cannot be accounted for by an existing network can always be accommodated by appropriate revisions. Similarly, any other theory of intensional phenomena can be reexpressed as a network theory. Thus, theories based on semantic features (see, e.g., McCloskey & Glucksberg, 1979; Schaeffer & Wallace, 1970; Smith et al., 1974) or meaning postulates (e.g., Fodor et al., 1975; Kintsch, 1974) can be readily translated into networks (see also Hollan, 1975). A major problem is precisely the lack of a general theory that restricts the class of potential networks to a readily testable subset. We have no quarrel with the formalism or notation of networks: A commitment to them is little more restrictive, and no more open to criticism, than is a commitment to a particular programming language such as LISP.

We have now described the main varieties of semantic networks, outlined their common underlying assumptions and their points of difference, and showed that the class of theories as a whole is insufficiently constrained. Our next task is to assess network theory in relation to the four major goals of semantic theory that we presented in the first section of the article.

Some Problematical Semantic Phenomena

The most striking design feature of semantic networks is that they have virtually nothing to say about extensional relations. In contrast, *model-theoretic* semantics, which derives from work in formal logic, provides a method of specifying the extensions of expressions within a model structure—typically, a set of numbers or some other abstract domain. The theorist posits the extensions of the basic lexical items in the model structure and formulates a set of semantic rules (which usually work in tandem with the syntactic rules of the language) for combining these extensions to form the extensions of expressions, and so on, all the way up to the truth value of the sentence. The method has been applied to natural language by Montague (1974) and others.

Some network theorists have recognized that networks lack any extensional machinery, and J. R. Anderson (1976) has even put forward a model-theoretic semantics for his ACT system that specifies the truth conditions of sentences in terms of the extensions of their constituents. Anderson's aim, unfortunately, is not to formulate a theory of how people represent extensional relations, but rather to use this semantics to determine the expressive power of the ACT system.

Most of the problems with current semantic networks arise from the neglect of extensions and from the assumption that the intensional relations between expressions can be analyzed independently from extensional matters. In this section, we are going to consider five phenomena: the relations among intensional relations, ambiguity, anomaly, instantiation, and inference. Any psychological theory of meaning should account for these phenomena; semantic networks contain mechanisms designed to do so, but nevertheless fail to deal with them adequately, a failing that also applies to theories based on semantic features or on meaning postulates.

Relations Among Intensional Relations

Although semantic networks are designed to represent intensional relations, they are not readily able to accommodate the full range of phenomena associated with them. As network theory acknowledges, there are different types of intensional relation (e.g., synonymy, antonymy, class inclusion, and partonymy). Moreover, native speakers can make judgments about types of intensional relation. They can judge, for example, that large is opposite to small, that slim is similar to thin, that silver contrasts with gold, and that a kitchen is a part of a house. Network theory can explain these judgments by using additional labelled links. For example, the words *large* and *small* might be joined by a link labelled as antonymous (see Glass, Holyoak, & Kiger, 1979). Alternatively, they might each be linked to a higher order node, such as size, with these links marked as inconsistent (Glass & Holyoak, 1974/1975). However, such links account for neither the patterns of similarity among the different types of intensional relation nor subjects' judgments about these patterns. Thus, synonymy (e.g., car–auto) is judged as more similar to class inclusion (car–vehicle) than it is to partonymy (car–engine). These differences in similarity are reliable across tasks and studies spanning several decades (Chaffin & Herrmann, 1981, 1984; Chaffin, Winston, & Herrmann, 1984; Moran, 1941; Perfetti, 1967; Riegel & Riegel, 1963; Warren, 1921).

Types of intensional relation can, of course, be analyzed and defined (see, e.g., Evens, Litowitz, Markowitz, Smith, & Werner, 1983), and subjects' judgments about the patterns of similarity can be explained in terms of their having access to some such underlying specifications (Chaffin, Russo, & Herrmann, 1981; Herrmann et al., 1979). Likewise, the latencies of their judgments about the intensional relation between a pair of words can be accounted for by the degree to which the actual relation between the pair conforms to the specification of the type of intensional relation. For example, judgments that two words are antonymous are faster if they are symmetrically opposed (e.g., large–small) than if they are asymmetrically opposed (e.g., large–tiny; Herrmann, Chaffin, & Daniel, 1984). Judgments that two words are synonymous are faster if they have the same sets of extensions (e.g., car–auto) than if the sets are merely similar (e.g., taxi–limousine; Herrmann, 1978; Herrmann, Papperman, & Armstrong, 1978). There are similar results with set inclusion (Loftus, 1973; Meyer, 1970) and part–whole judgments (Chaffin et al., 1984). Because current networks do not represent the underlying specifications of intensional relations, which are treated as unanalyzed labels on links, they cannot account for judgments about types of intensional relation.

There are various ways in which the required machinery might be introduced into networks. The most natural method would be to represent the underlying semantics of the different types of intensional relation within the network itself. It would be necessary to formulate definitions of synonymy, antonymy, and the other types of relation in terms of more primitive notions and then to construct network representations of these definitions. For instance, one component of antonymy is that the two terms should denote opposite values of the same property; *tall* and *small* are not antonyms because the former applies to height and the latter, to size. Likewise, if two terms deal with an underlying continuum, then they should be symmetric about the midpoint; *hot* and *cold* are antonyms, whereas *hot* and *tepid* are not. Similar intensional relations would obviously have similar representations in the network. This pattern would account for subjects' judgments of similarity. The degree to which the links between a pair of words conformed to the definition of a given intensional relation would likewise account for the latencies of such judgments (Herrmann & Chaffin, 1984). What is unclear, however, is the nature of the links required to define the various types of intensional relation and the nature of the processes for comparing definitions with the actual links between pairs of words. Given the power of networks, we have no doubt that such machinery could be developed, but even this addition leaves further phenomena unexplained.

In particular, there are analogous problems that arise with relations that hold, not between the senses of words, but between what they are actually used to refer to—their extensions. For example, anyone who has ever packed goods at a supermarket is familiar with the fact that some things are more easily squashed than others. Tomatoes are more squashable than potatoes. Semantic networks are supposed to be appropriate representations for all concepts, but no feasible network could contain a label between tomatoes and potatoes that expresses this relation (Chaffin & Herrmann, 1984). Moreover, if one were to ignore the exponential growth in the number of links that would result from adding such links, then there is still a problem. If potatoes are cooked and mashed, or if tomatoes are frozen solid, then it is no longer the case that tomatoes are more squashable than potatoes. Judgments about such relations do indeed require access to the extensions of expressions.

Other relations cause problems because they are ad hoc and more abstract or complex. They include Boolean functions of simple relations, like the set of all pairs, X and Y, such that X is larger and faster then Y, and quantified relations like the set of all pairs of sets, X and Y, such that some members of X are taller than all members of Y. The latter has quantifiers that range over the individuals in two sets (e.g., the relation holds between men and women, adults and children, animals and human beings, but it does not hold between the converse pairs). Because only a fraction of all possible Boolean and quantified relations would be directly represented in a network, the evaluation of a relation would require in most cases a computation that compared the range of values associated with the two terms. For example, the verification of "An airplane is larger and faster than a squirrel" could not be done by a search through the network because there would be no direct links labelled "larger than" and "faster than" connecting the two terms. It would therefore be necessary to compare the respective sizes and speeds associated with airplane and squirrel and to determine whether the differences between them satisfied the sense of "larger and faster than."

Of course, the existence of a mechanism for evaluating the definitions of intensional relations does not rule out the possibility that some intensional relations are directly represented in a network (or some other form of intensional representation, such as a set of semantic features or meaning postulates). What relations might be encoded in this way? A plausible answer is those relations that are directly learned in a metalinguistic fashion. When a child learns that little is the opposite of big, or that large is similar in meaning to big, then these metalinguistic facts could be directly represented by labelled links between the representations of the appropriate words. Most intensional relations, however, are not directly acquired for all the terms that satisfy them (e.g., children do not learn the set of opposites exhaustively). Likewise, many intensional relations themselves are never explicitly learned. It follows that the direct representation of intensional relations is likely to play only a small part in semantic judgments.

Ambiguity

Comprehensive network theories are designed to represent the meanings of sentences. They therefore have to be able to cope with the resolution of ambiguities. Many words are in fact ambiguous, and the more frequently used words are more likely to be ambiguous (see Miller, 1951). Speakers and listeners seldom notice these ambiguities unless the sentence as a whole is ambiguous, and hence the mechanisms for resolving lexical ambiguity must be highly efficient. Semantic networks cope with ambiguous words by using so-called "selectional restrictions" to determine their appropriate interpretations. This idea derives from a standard procedure in lexicography: Dictionaries relate each different sense of a word to the meanings of other words that can occur in construction with it. The meanings of these other words accordingly select the appropriate sense as shown in the following example:

plant (v.t.): to set or sow (seeds or plants) in the earth;
to establish (animals) in a new locality,

where the parenthetical items in the definition are the selectional restrictions. Katz and Fodor (1963) developed this approach into a more formal theory, and Quillian (1968) and his intellectual descendants adopted it as well. An ambiguous word is disambiguated by establishing the words that occur in construction with it in the same sentence and checking which restrictions stated in its network definition their meanings satisfy. Hence, the sentence, "He planted his herd on the island," cannot mean that he sowed the earth with animals because this sense of the verb is restricted to seeds or plants.

Unfortunately, selectional restrictions have only a limited explanatory value. What is needed in disambiguating a word is access to the extensions of other words in the sentence. This point becomes clear from considering how the verb in the following sentence is interpreted: "He planted them on the island." If, for instance, the pronoun "them" refers to seeds, then the verb will be taken to mean that the seeds were sown. However, the pronoun does not have seeds as one of its meanings; rather, it is the extension of the term. At the very least, therefore, the apparatus of selectional restrictions embodied in semantic networks must be modified so that it has access to extensions.

There is a still more serious problem. It is impossible to state complete and consistent selectional restrictions on the meanings of many ambiguous words. Consider, for instance, the transitive verb "to lift," which has a number of different meanings, including: to move (an object) from a lower to a higher position; to steal or plagiarize (an idea); and to put an end to (a siege or blockade or embargo). What restrictions does the first sense place on the subject of a sentence of the form "*X* lifted *Y*"? As a first approximation, one might suppose that *X* should be human, animal, or a machine. However, this hypothesis fails to accommodate the following sentence: "The wind lifted the leaves." It seems the selectional restriction should be: capable of exerting a physical force, and that every item in the network denoting a human, animal, machine, or entity such as the wind should be linked to a node representing this meaning. However, such a scheme will still not work properly. Further machinery is needed to cope with the fact that the particular entity referred to by the subject of the sentence may, or may not, be capable of exerting a force depending on the circumstances to which the assertion refers. Thus, for example, although human beings are generally capable of exerting a force, dead or paralyzed human beings are not.

A sensible construal of the theory of selectional restrictions is that it captures those inferences that are made so often that they have become assumptions about how words are to be interpreted by default (see Miller & Johnson-Laird, 1976, p. 701). Hence, unless there is evidence to the contrary, the subject of "lift," if it is to be interpreted as meaning move upward, should be animate or machine-like. In general, however, lexical ambiguities can be resolved only by inferences based on knowledge about the specific entities referred to in the sentence.

Anomaly

Semantic anomaly arises whenever a sentence cannot be given a meaningful interpretation. For example, the sentence, "The bowl of soup ate the ham sandwich" on any feasible theory of selectional restrictions ought to be anomalous, because only animate entities can eat sandwiches. The sentence might, of course, have the metaphorical interpretation that the ham sandwich fell into the bowl of soup and was engulfed by it; and this interpretation might be explained by mechanisms that come into play only when the system detects a semantic anomaly. As Bolinger (1965) wrote, "It is characteristic of natural language that no word is ever limited to its enumerable senses, but carries with it the qualification, 'something like'." However, there is quite another interpretation of the sentence above that is not metaphorical. Nunberg (1978) pointed out that such noun phrases as "the bowl of soup" and "the ham sandwich" could be readily used by a waiter as a way of referring to his customers in terms of what they have ordered. A theory that assumes that a set of literal meanings is first assigned to a sentence and that context then eliminates inappropriate meanings runs into grave difficulties with such sentences. Once again, the proper treatment of this intensional phenomenon calls for semantic networks to have access to extensions.

Instantiation

The linguistic context in which an unambiguous word occurs can narrow down its interpretation. This phenomenon of "instantiation" has been demonstrated experimentally by R. C. Anderson and his colleagues. They presented subjects with a sentence such as "The fish attacked the swimmer," and later the subjects had to recall the sentence (see R. C. Anderson et al., 1976). A more specific and likely term such as *shark* was a better recall cue than the original general term, *fish*. Analogous phenomena occur with the interpretation of verbs (Garnham, 1979), and context can also render one aspect of the meaning of a word more salient than another (Tabossi & Johnson-Laird, 1980).

Halff, Ortony, and Anderson (1976) accounted for instantiation by arguing that words do not have a few qualitatively distinct meanings, but rather a whole family of potential meanings (see also Putnam, 1975; Weinreich, 1966). When a word occurs in a sentence, the linguistic context acts to instantiate a specific member of the family of meanings. This argument is mildly embarrassing for network theories, because it implies that each word should be associated with many more concept nodes than theorists normally envisaged. However, there is an alternative explanation of instantiation that is both more convincing and more embarrassing for network theories. The sentence, "It attacked the swimmer," might well be better recalled given the cue *shark* rather than the original term, *it*. No one would seriously argue that "it" is polysemous; it has a single meaning, which enables it to refer to an infinite number of different entities. Hence, what is instantiated by linguistic context is not a particular sense but a particular referent. There is no need to invoke vast sets of meanings for words, but there is a need for access to extensions.

Inference

One function of semantic networks is to enable inferences to be drawn from verbal assertions. For example, the deduction, "Fido is a poodle; therefore, Fido is a dog" depends on the unstated premise that all poodles are dogs. The inference can be made by traversing first the link that establishes that Fido is a member of the set of poodles and then the link that establishes that the set of poodles is included in the set of dogs. This method of making inferences is formal in that one expression is derived from others by virtue of a configuration of symbols and without regard to their extensions or truth conditions. Indeed, J. R. Anderson argued for the network representation of HAM in part because it suggested this sort of inferential mechanism (cf. J. R. Anderson, 1976, p. 41). However, there are valid inferences that cannot be made in a formal way, because the interpretation of terms depends on the situation referred to rather than on the senses of the expressions. For example, inferences of the form, *A* is on *B*'s right; *B* is on *C*'s right; therefore, *A* is on *C*'s right, are valid provided that *A*, *B*, and *C* are seated down one side of rectangular table and facing in the same direction. The inference may be invalid if the relevant individuals are seated around a circular table. It is of little use to claim that the expression "on *X*'s right" has two meanings—one transitive and one intransitive—because with a seating arrangement of an appropriate radius transitivity can extend over some arbitrarily finite number of individuals (Johnson-Laird, 1981). These vagaries of deductive inference require an inferential system that has access to a representation of the situation to which reference is made.

Extensional Models and Redundancy of Intensional Relations

Symbolic Fallacy

In the previous section, we outlined some problems that arise if intensional relations are supposed to be interpreted independently from extensions. Semantic networks are, of course, based on that assumption. Indeed, theorists have assumed that the process of translating natural language into a network representation amounts to comprehension. They have overlooked the fact that such a representation is strictly meaningless unless in principle it can be connected to a conception of the external world; and, of course, network theory does not establish any such connection. In this respect, semantic networks render theorists particularly vulnerable to what we call the "symbolic" fallacy—the assumption that the mere translation of sentences into symbols constitutes a useful account of their meaning (see also Lewis, 1972). For example, Kintsch and van Dijk (1978) proposed a propositional representation for sentential meanings, and they were sensitive to the potential emptiness of the procedure. When they translated "The students complain" into *(complain, student)*, they pointed out:

> *Complain* is merely the name of a knowledge complex that specifies the normal uses of this concept, including antecedents and consequences. Inferences can be derived from this knowledge complex, for example, that the students were probably unhappy or angry or that they complained to someone, with the text supplying the professor for that role. Once a concept like *complain* is elaborated in that way, the semantic notation is anything but vacuous. (Kintsch & Van Dijk, 1978, p. 378)

Unfortunately, if the knowledge complex merely consists of other symbolic expressions, then there is no escape from the maze of symbols into the world: The theory has succumbed to the symbolic fallacy. The fallacy is analogous to the hackneyed science fiction idea that the aliens learned the languages of the Earth's inhabitants from radio transmissions. One might (just) learn that "zug brochna" follows from "gek brochna" in this way, but no matter how complex the inferences that can be acquired, the system will never amount to a proper semantics. It connects expressions to expressions perhaps, but it does not connect them to their extensions. This issue is precisely the one that cropped up in different guises with the problematical phenomena of the previous section: They could not be handled by networks, which lack representations of the extensions of expressions. Of course, network theories are not alone in committing the symbolic fallacy: A reluctance to deal with reference is evident in theories based on semantic features or on meaning postulates.

Models of Extensions

Is it possible to remedy network theory so that it deals with extensions and thus copes with the puzzling aspects of intension relations? The answer is both yes and no. Any putative theory can be reexpressed in the form of a network system because, as we mentioned earlier, networks in general have the power of Universal Turing machines. In contrast, as we show, the required solution appears to call for structures that are foreign to network theory, and so the revised networks will not look at all like the sorts of theory that we have considered in this review.

A fruitful way in which to approach the problem of extensions is to develop a psychologically oriented model-theoretic semantics based on the assumption that sentences relate to the world by virtue of mental models of states of affairs (Johnson-Laird, 1983). Human beings can construct models of the world on the basis of perception, memory, and imagination; they can also construct rudimentary models of the world on the basis of verbal descriptions; and they can sometimes evaluate an expression by comparing its linguistically derived model with a perceptual or conceptual one. In many cases, of course, the process of comparison may be difficult: If the truth conditions concern social relations or states of affairs that have no direct perceptual correlate in the world, then the relevant evidence may be at best indirect. In other cases, the process of comparison may be impossible: If the truth conditions of the assertion are incomplete or indeterminate, then no evidence may suffice to verify the assertion. Nevertheless, the fact that human beings can in principle relate language to models of the world provides the foundation of semantics.

A semantic theory for mental models requires two levels of representation. First, an utterance has to be translated into an intensional representation—a representation of the proposition that the utterance expresses. Second, the propositional representation may be translated into an extensional representation—a mental model of the particular state of affairs characterized by the utterance. The two levels are required in order to account for a number of inferential and interpretative phenomena, and they are borne out by experimental evidence (e.g., Mani & Johnson-Laird, 1982; Ehrlich & Johnson-Laird, 1982). Most important, they allow for the recursive revision of the model if, as in the case of an indeterminate description, it should embody some assumption that subsequently turns out to be erroneous. In this way, a single specific model can represent a proposition: The system sets up a representative sample from the set of possible models of the discourse, and if the sample fails to match a subsequent assertion, it can be revised (within the constraints of human memory) so as to be consistent with the discourse as a whole. This approach solves the problem of how a single model of, for example, a cat sitting on a mat can represent the indefinite number of different states of affairs that are truthfully described by a sentence to that effect. The procedures that revise the model, however, must have access to an independent record of the sentence, and the initial propositional representation serves this function.

The propositional representation of a sentence may well take the form of a semantic network, but plainly there is a need for a further extensional representation, or mental model, of the state of affairs described by the sentence. Mental models, unlike semantic networks, have a structure that corresponds to the perceived or conceived structure of that state of affairs. Thus, for example, a mental model of an assertion such as "Every person owns a car" contains an arbitrary number of tokens corresponding to persons, an arbitrary number of tokens corresponding to cars, and relations between these two sets of tokens designating ownership:

Person 1 → Car 1

Person 2 → Car 2

Person 3 → Car 3

(Car 4)

The parenthetical item represents the fact that there may be cars that are not owned by anyone. It is important to note that the structural constraint on mental models (i.e., that their structure corresponds to the perceived or conceived structure of states of affairs) renders them quite distinct from conventional semantic networks. Thus, the network representation of the same sentence, which is shown in Figure 7, calls for four partitions, eight nodes, and eight links—a structure that is remote from the structure of an actual state of affairs in which every person owns a car, where indeed there will be two sets with a relation from all the members of one set to members of the other.

If mental models are to be constructed from sentences, the meanings of words have to specify the conditions that must be satisfied in the model for the word to apply truthfully to it. The lexical semantics must therefore capture the contribution that the word makes to the truth conditions of any assertion in which it occurs. Thus, for example, the semantics for the expression "on the right of" specifies the direction to be scanned in order to form a mental model of such assertions as "*A* is on the right of *B*." The interpretative system locates one entity within a spatial array and then scans in an appropriate direction with respect to the observer's viewpoint in order to locate the other. The direction can be specified straightforwardly (e.g., using iteration on Cartesian coordinates), but obviously such a specification calls for concepts that may well be ineffable in the language under analysis. Indeed, one cannot define the truth conditions for the sentence "*A* is on the right of *B*" in ordinary language.

The conjecture that subjects construct mental models of the events described in discourse goes a considerable way toward explaining the resolution of ambiguities, anomalies, instantiation, the vagaries of inference, and the well-known findings of Bransford and his colleagues that subjects often go beyond what is linguistically given in interpreting discourse (Bransford & McCarrell, 1975).

Specifications of Intensional Relations May Be Redundant

If a computer program is based on the theory of mental models (see Johnson-Laird, 1983), a striking consequence emerges: It can work without relying on any representation of intensional relations of the sort captured in semantic networks. Thus, the program can make transitive inferences without recourse to rules of transitive inference. For example, the truth conditions for the assertion "*A* is on the right of *B*" enable an appropriate model to be constructed: *B A*. Likewise, the truth conditions for the further assertion "*B* is on the right of *C*" enable the program to extend the model appropriately to: *C B A*. The program contains a verification routine that is elicited whenever an assertion makes no reference to any new entities. The same truth conditions for "on the right of" therefore suffice to return the value, "true", when a third assertion is made: "*A* is on the right of *C*." At this point, of course, no valid deduction has been made. However, the program also contains a recursive procedure that searches for an alternative model of the premises that would render the current assertion false. There is no such alternative in this case, and (the program concludes that) the deduction is valid. In summary, the program makes a transitive inference without needing any rule to the effect that "on the right of" is transitive: Transitivity is an emergent property from the truth conditions of the relation.

The simplest way in which to state the redundancy of intensional relations is in terms of the following example. If you know what it is for something to be a canary and you know what it is for something to be a bird, then you do not need to know explicitly that canaries are birds, because this relation will be a consequence of your knowledge. Of course, the fact that information is redundant does not imply that it is not mentally represented.

We have reached the same conclusion that we arrived at in considering relations among intensional relations. Some intensional relations will be mentally represented, others will not. Those that are not mentally represented can always be recovered on the basis of processes that examine the truth conditions of terms. Which relations are likely to be redundantly specified? The answer, once again, is those that are directly learned. Hence, if you learn that canaries are birds, then this relation will be mentally encoded even though you may also learn about the extensions of the two terms. In contrast, there are terms for which you initially acquire a direct, although ineffable, knowledge of truth conditions. These terms are likely to include what Rosch (1976) called "basic-level" words, that is, words that denote entities that have attributes and functions maximally in common.

Conclusions

How do semantic networks fare in relation to the four goals for a psychological theory of meaning that we outlined earlier? They specify a form for the mental representation of meaning (the first goal), although this form is remote both from the immediate structure of the sentence and from the conceived structure of the world. They also account for some intensional properties and relations (the second goal), but they run into difficulties with certain cases of relations between words. They are not designed to specify the extensions of words or expressions (the third goal). In addition, their various accounts of inference (the fourth goal), whether by decomposition or by rules, are not entirely adequate for coping with the vagaries of validity.

We can state our misgivings about semantic networks by way of an example. One of the things that you know about poodles is that they are dogs, and this sort of information can be represented in a semantic network. However, you also have some knowledge about what it is for something to be a poodle; you have a concept of what poodles are, and this knowledge enables you to identify poodles, to establish complex intensional relations involving them, and to verify assertions about them. Of course, no one can definitively classify any entity as either a poodle or not a poodle: There are probably no necessary and sufficient conditions for poodlehood (see Putnam, 1975), but without some knowledge of what determines the extension of a term you can hardly be said to have grasped its meaning. This sort of knowledge is not represented in semantic networks, and there is no immediate way in which it could be represented because networks lack connections to representations of the world. They only provide connections between words. Unfortunately, they cannot even give a complete explanation of intensional phenomena, because a proper account of ambiguity, anomaly, instantiation, and inference turns out to depend on access to extensional representations. Because, as we noted earlier, network theory is a notational variant both of models based on semantic features (e.g., McCloskey & Glucksberg, 1979; Schaeffer & Wallace, 1970; Smith et al., 1974) and of models based on meaning postulates (e.g., Fodor et al., 1975; Kintsch, 1974), the shortcomings apply equally to these theories.

When a theory is constructed to handle extensions, then it becomes clear that it is no longer strictly necessary to represent the intensional connections between words. Any such theory that accommodates extensions could undoubtedly be reexpressed within the formalism of a network system. But this translation would be very different from current network theory. The new theory would be serving an extensional function, not addressed by previous network theories, and the structures of its network would be radically different because they would now correspond to the structures with which human beings conceive states of affairs.

Despite these strong qualifications on network theory, it remains highly likely that certain intensional relations are directly learned as such, and that they may accordingly be represented by associative relations between representations of words. Semantic networks may therefore still have a role, albeit a reduced one, in accounting for the mental representation of the meanings of words. They may also have a major role as a hypothesis about the structure of the initial encoding of sentences especially if they are supplemented by some form of extensional representation. The moral of this article is simple. Semantic networks can do many things, but they cannot do everything that a psychological theory of meaning ought to be able to do: The meanings of words can only be properly connected to each other if they are properly connected to the world.

References

Anderson, J. R. (1976). *Language, memory, and thought.* Hillsdale, NJ: Erlbaum.

Anderson, J. R., & Bower, G. H. (1973). *Human associative memory.* Washington, DC: Winston.

Anderson, R. C., Pichert, J. W., Goetz, E. T., Schallert, D. L., Stevens, K. V., & Trollip, S. R. (1976). Instantiation of general terms. *Journal of Verbal Learning and Verbal Behavior, 15,* 667–679.

Bierwisch, M. (1970). Semantics. In J. Lyons (Ed.), *New horizons in linguistics* (pp. 166–184). Harmondsworth, Middlesex, England: Penguin.

Bilger, H. (1979). On the notion of irreducible relations. In G. Bracchi & G. M. Nijssen (Eds.), *Data base architecture* (pp. 277–296). Amsterdam: North-Holland.

Bobrow, D. G., & Collins, A. M. (Eds.). (1975). *Representation and understanding: Studies in cognitive science.* New York: Academic Press.

Bolinger, D. (1965). The atomization of meaning. *Language, 41,* 555–573.

Brachman, R. J. (1979). On the epistemological status of semantic networks. In N. V. Findler (Ed.), *Associative networks: The representation and use of knowledge by computers* (pp. 3–50). New York: Academic Press.

Bransford, J. D., & McCarrell, N. S. (1975). A sketch of a cognitive approach to comprehension: Some thoughts about understanding what it means to comprehend. In W. B. Weimar & D. S. Palermo (Eds.), *Cognition and the symbolic processes* (pp. 189–230). Hillsdale, NJ: Erlbaum.

Chaffin, R., & Herrmann, D. J. (1981). Semantic relationships and the generality of semantic-memory models. *Bulletin of the Psychonomic Society, 12,* 69–72.

Chaffin, R., & Herrmann, D. J. (1984). The diversity and similarity of semantic relations. *Memory and Cognition, 12,* 134–141.

Chaffin, R., Russo, A., & Herrmann, D. J. (1981). An effect of relationship similarity on categorization latency. *Journal of General Psychology, 104,* 305–306.

Chaffin, R., Winston, M., & Herrmann, D. J. (1984). *The semantics of part–whole relations: An empirical taxonomy.* Unpublished manuscript.

Chomsky, N. (1965). *Aspects of the theory of syntax.* Cambridge, MA: MIT Press.

Clark, H. H., & Clark, E. V. (1977). *Psychology and language: An introduction to psycholinguistics.* New York: Harcourt Brace Jovanovich.

Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review, 82,* 407–428.

Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior, 8,* 240–247.

Collins, A. M., & Quillian, M. R. (1972). How to make a language user. In E. Tulving & W. Donaldson (Eds.), *Organization and memory* (pp. 310–349). New York: Academic Press.

Conrad, C. (1972). Cognitive economy in semantic memory. *Journal of Experimental Psychology, 92,* 149–154.

Deese, J. (1962). On the structure of associative meaning. *Psychological Review, 69,* 161–175.

Deese, J. (1965). *The structure of associations in language and thought.* Baltimore, MD: The Johns Hopkins University Press.

Ehrlich, K., & Johnson-Laird, P. N. (1982). Spatial descriptions and referential continuity. *Journal of Verbal Learning and Verbal Behavior, 21,* 296–306.

Evens, M. W., Litowitz, D. E., Markowitz, J. A., Smith, R. N., & Werner, D. (1983). Lexical-semantic relations: A comparative survey. *Current inquiry into language and linguistics* (Report No. 34). Carbondale, IL and Edmundton, Alberta, Canada: Linguistic Research Inc.

Fillmore, C. J. (1968). The case for case. In E. Bach & R. T. Harms (Eds.), *Universals in linguistic theory* (pp. 1–88). New York: Holt, Rinehart & Winston.

Findler, N. V. (Ed.). (1979). *Associative networks: Representation and use of knowledge by computers.* New York: Academic Press.

Fodor, J. D., Fodor, J. A., & Garrett, M. F. (1975). The psychological unreality of semantic representations. *Linguistic Inquiry, 4,* 515–531.

Frege, G. (1892). Über sinn und bedeutung [On sense and reference]. *Zeitschrift für Philosophie und philosophische Kritik, 100,* 25–50. In P. T. Geach & M. Black (Eds. and Trans.), *Translations from the philosophical writings of Gottlob Frege.* Oxford: Blackwell, 1952.

Garnham, A. (1979). Instantiation of verbs. *Quarterly Journal of Experimental Psychology, 31,* 207–214.

Glass, A. L., & Holyoak, K. J. (1974/1975). Alternative conceptions of semantic memory. *Cognition, 3,* 313–339.

Glass, A. L., Holyoak, K. J., & Kiger, J. I. (1979). Role of antonymy relations in semantic judgments. *Journal of Experimental Psychology: Human Learning and Memory, 5,* 598–606.

Halff, H. M., Ortony, A., & Anderson, R. C. (1976). A context-sensitive representation of word meanings. *Memory and Cognition, 4,* 378–383.

Hendrix, G. G. (1979). Encoding knowledge in partitioned networks. In N. V. Findler (Ed.), *Associative networks: Representation and use of knowledge by computers* (pp. 51–92). New York: Academic Press.

Herrmann, D. J. (1978). An old problem for the new psychosemantics: Synonymity. *Psychological Bulletin, 85,* 490–512.

Herrmann, D. J., & Chaffin, R. (1984). *Definitions of sense relations.* Unpublished manuscript.

Herrmann, D. J., Chaffin, R., & Daniel, M. (1984). *Three attributes of antonymy: Effects of relation attributes on antonym comprehension.* Manuscript submitted for publication.

Herrmann, D. J., Chaffin, R. J. S., Conti, G., Peters, D., & Robbins, P. H. (1979). Comprehension of antonymy and the generality of categorization models. *Journal of Experimental Psychology: Human Learning and Memory, 5,* 585–597.

Herrmann, D. J., Papperman, T. J., & Armstrong, A. D. (1978). Synonym comprehension and the generality of categorization models. *Memory and Cognition, 6,* 150–155.

Hollan, J. D. (1975). Features and semantic memory: Set-theoretic or attribute model? *Psychological Review, 82,* 154–155.

Holyoak, K. J., & Glass, A. L. (1975). The role of contradictions and counterexamples in the rejection of false sentences. *Journal of Verbal Learning and Verbal Behavior, 14,* 215–239.

Humphrey, G. (1951). *Thinking: An introduction to its experimental psychology.* London: Methuen.

Johnson-Laird, P. N. (1970). The interpretation of quantified sentences. In G. B. Flores d'Arcais & W. J. M. Levelt (Eds.), *Advances in psycholinguistics* (pp. 347–372). Amsterdam: North-Holland.

Johnson-Laird, P. N. (1974). Experimental psycholinguistics. *Annual Review of Psychology, 25,* 135–160.

Johnson-Laird, P. N. (1981). Mental models of meaning. In A. K. Joshi, B. L. Webber, & I. A. Sag (Eds.), *Elements of discourse understanding* (pp. 106–126). Cambridge, England: Cambridge University Press.

Johnson-Laird, P. N. (1983). *Mental models: Towards a cognitive science of language, inference, and consciousness.* Cambridge, England: Cambridge University Press.

Katz, J. J. (1972). *Semantic theory.* New York: Harper & Row.

Katz, J. J., & Fodor, J. A. (1963). The structure of a semantic theory. *Language, 39,* 170–210.

Kintsch, W. (1974). *The representation of meaning in memory.* Hillsdale, NJ: Erlbaum.

Kintsch, W., & van Dijk, T. A. (1978). Towards a model of text comprehension and reproduction. *Psychological Review, 85,* 363–394.

Kiss, G. R. (1967). Networks as models of word storage. In N. L. Collins & D. Michie (Eds.), *Machine intelligence, Vol. I* (pp. 155–168). Edinburgh: Oliver & Boyd.

Landauer, T. K., & Freedman, J. L. (1968). Information retrieval from long-term memory: Category size and recognition time. *Journal of Verbal Learning and Verbal Behavior, 7,* 291–295.

Lewis, D. K. (1972). General semantics. In D. Davidson & G. Harman (Eds.), *Semantics of natural language* (pp. 169–218). Dordrecht, Holland: Reidel.

Lindsay, P. H., & Norman, D. A. (1977). *Human infor-*

mation processing: An introduction to psychology (2nd ed.). New York: Academic Press.

Loftus, E. F. (1973). Category dominance, instance dominance, and categorization time. *Journal of Experimental Psychology, 97,* 70–74.

Lyons, J. (1977). *Semantics* (Vol. 1). Cambridge, England: Cambridge University Press.

Mani, K., & Johnson-Laird, P. N. (1982). The mental representation of spatial descriptions. *Memory and Cognition, 10,* 181–187.

McCloskey, M., & Glucksberg, S. (1979). Decision processes in verifying category membership statements: Implications for models of semantic memory. *Cognitive Psychology, 11,* 1–37.

Meyer, D. E. (1970). On the representation and retrieval of stored semantic information. *Cognitive Psychology, 21,* 242–300.

Miller, G. A. (1951). *Language and communication.* New York: McGraw-Hill.

Miller, G. A., & Johnson-Laird, P. N. (1976). *Language and perception.* Cambridge, MA: Harvard University Press.

Montague, R. (1974). *Formal philosophy: Selected papers.* New Haven, CT: Yale University Press.

Moran, L. (1941). Generality of word association response sets. *Psychological Monographs, 30,* (4, Whole No. 612).

Norman, D. A., Rumelhart, D. E., & the LNR Research Group (Eds.). (1975). *Explorations in cognition.* San Francisco: Freeman.

Nunberg, G. (1978). *The pragmatics of reference.* Bloomington, IN: Indiana University Linguistics Club.

Perfetti, C. A. (1967). A study of denotative similarity with restricted word associations. *Journal of Verbal Learning and Verbal Behavior, 6,* 788–795.

Putnam, H. (1975). The meaning of "meaning." In K. Gunderson (Ed.), *Language, mind and knowledge: Minnesota studies in the philosophy of science* (Vol. 7, pp. 131–193). Minneapolis: University of Minnesota Press.

Quillian, M. R. (1968). Semantic memory. In M. L. Minsky (Ed.), *Semantic information processing* (pp. 227–259). Cambridge, MA: MIT Press.

Raphael, B. (1968). SIR: A computer program for semantic information retrieval. In M. L. Minsky (Ed.), *Semantic information processing* (pp. 33–110). Cambridge, MA: MIT Press.

Riegel, K. F., & Riegel, R. M. (1963). An investigation into denotative aspects of word meaning. *Language and Speech, 6,* 5–21.

Rips, L. J., Shoben, E. J., & Smith, E. E. (1973). Semantic distance and the verification of semantic relations. *Journal of Verbal Learning and Verbal Behavior, 12,* 1–20.

Rogers, H. (1967). *Theory of recursive functions and effective computability.* New York: McGraw-Hill.

Rosch, E. (1973). On the internal structure of perceptual and semantic categories. In T. M. Moore (Ed.), *Cognitive development and the acquisition of language* (pp. 111–144). New York: Academic Press.

Rosch, E. (1976). Classification of real-world objects: Origins and representations in cognition. In S. Ehrlich & E. Tulving (Eds.), *La memoire semantique* (pp. 242–250). Paris: Bulletin de Psychologie. (Reprinted in P. N. Johnson-Laird & P. C. Wason (Eds.), *Thinking: Readings in cognitive science* (pp. 212–222). Cambridge, England: Cambridge University Press, 1976).

Rumelhart, D. E., & Levin, J. A. (1975). A language comprehension system. In D. A. Norman, D. E. Rumelhart, & the LNR Research Group (Eds.), *Exploration in cognition* (pp. 179–208). San Francisco: Freeman.

Rumelhart, D. E., Lindsay, P. H., & Norman, D. A. (1972). A process model for long-term memory. In E. Tulving & W. Donaldson (Eds.), *Organization and memory* (pp. 198–246). New York: Academic Press.

Schaeffer, B., & Wallace, R. (1970). The comparison of word meanings. *Journal of Experimental Psychology, 86,* 144–152.

Shoben, E. J., Wescourt, K. T., & Smith, E. E. (1978). Sentence verification, sentence recognition, and the semantic-episodic distinction. *Journal of Experimental Psychology: Human Learning and Memory, 4,* 304–317.

Smith, E. E., Shoben, E. J., & Rips, L. J. (1974). Structure and process in semantic memory: A featural model for semantic decisions. *Psychological Review, 81,* 214–241.

Steinberg, D. D. (1970). Analyticity, amphigory, and the semantic interpretation of sentences. *Journal of Verbal Learning and Verbal Behavior, 1,* 37–51.

Tabossi, P., & Johnson-Laird, P. N. (1980). Linguistic context and the priming of semantic information. *Quarterly Journal of Experimental Psychology, 32,* 595–603.

Walpole, H. R. (1941). *Semantics: The nature of words and their meanings.* New York: Norton.

Warren, H. C. (1921). *A history of association in psychology.* New York: Scribner's.

Weinreich, U. (1966). Explorations in semantic theory. In T. A. Sebeok (Ed.), *Current trends in linguistics, Vol. 3* (pp. 395–477). The Hague, The Netherlands: Mouton.

Woods, W. A. (1975). What's in a link: Foundations for semantic networks. In D. G. Bobrow & A. M. Collins (Eds.), *Representation and understanding: Studies in cognitive science* (pp. 35–82). New York: Academic Press.

Received September 9, 1983
Revision received April 5, 1984 ■