9 / Anthony S. Maida and Stuart C. Shapiro
Intensional Concepts in Propositional Semantic Networks

This paper proposes a particular semantic network formalism apparently capable of handling three important problems having to do primarily with natural language processing. One of these—the "telephone number problem"—is presented in McCarthy's "First Order Theories" paper [Chapter 31] in this volume. But more interestingly, the authors take and defend a very strong position with respect to the meanings of the nodes in their networks. Partially in response to [Woods, Chapter 11], but going substantially further, they interpret all information in the network *intensionally*—that is, as representing concepts and not items or sets in the world. The claimed value of this strict intensional stance is the ability of the formalism to handle tricky cases of sentences about beliefs. Since their nets are taken to model "the belief structure of a thinking, reasoning, language using being," Maida and Shapiro attempt to take opaque contexts (where, for example, equals cannot be substituted for equals) as the norm. While related in some ways to other formalisms discussed in this book, this creates an interesting contrast to KL-ONE [Brachman, Chapter 10], KRYPTON [Brachman *et al.*, Chapter 24], and KRL [Bobrow and Winograd, Chapter 13]. For example, the authors claim that their nets have no structural or definitional information, only assertions (which they claim is consistent with network models in the psychological literature). Further, the authors explain that in their model there can be no semantic primitives of the sort discussed in [Schank and Rieger, Chapter 7], and indeed draw upon the work of Quillian (see [Quillian, Chapter 6]) in defense of the non-compositional nature of their networks. There is, however, a bit of confusion in this paper, for example, when the authors get involved in discussions of "computational overhead" or in the direct representation of natural language questions, but the paper is valuable nevertheless. In particular, it is very well connected to the literature, and provides some good perspective on a number of other papers in this volume.

Appeared in *Cognitive Science* 6(4), 1982, 291-330.

Intensional Concepts in Propositional Semantic Networks*

ANTHONY S. MAIDA
Brown University
STUART C. SHAPIRO
State University of New York at Buffalo

An integrated statement is made concerning the semantic status of nodes in a propositional semantic network, claiming that such nodes represent only intensions. Within the network, the only reference to extensionality is via a mechanism to assert that two intensions have the same extension in some world. This framework is employed in three application problems to illustrate the nature of its solutions.

The formalism used here utilizes only assertional information and no structural, or definitional, information. This restriction corresponds to many of the psychologically motivated network models. Some of the psychological implications of network processes called node *merging* and node *splitting* are discussed. Additionally, it is pointed out that both our networks and the psychologically based networks are prone to memory confusions about knowing unless augmented by domain-specific inference processes, or by structural information.

INTRODUCTION

In this paper, we discuss a particular kind of semantic network. It is a representation of knowledge consisting of nodes and labeled, directed arcs in which the following conditions hold (cf. Shapiro, 1971): 1) each node represents a unique concept; 2) each concept represented in the network is represented by a node; 3) each concept represented in the network is represented by a unique node (the *Uniqueness Principle*); 4) arcs represent non-conceptual binary relations between nodes; 5) the knowledge represented about each concept is represented by the structure of the entire network connected to the node representing the concept. The term *propositional* semantic network is sometimes used to distinguish those semantic networks in which *every* assertion that can be stored or accessed in the network is considered a concept, and therefore represented by a node, from those networks in which some such assertions are represented by arcs (most notably set membership and subset relationships, cf. Hendrix, 1979, p. 54; or statements without handle nodes, cf. Fahlman, 1979, p. 112). This paper is concerned with propositional semantic networks so understood. Conceivably, conceptual dependency networks (cf. Schank, 1975) could be classified as propositional semantic networks if their syntax were explicated. We are using the term proposition because propositions are the intensions of sentences. The definition also allows for the use of nodes corresponding to functional individuals.

We will more closely examine the Uniqueness Principle in order to help understand the semantics of semantic networks. In doing this, we follow the line of research exemplified by Woods (1975) and Brachman (1977, 1979). The major point will be that nodes represent only intensions and not extensions, e.g., individual concepts rather than referents, and propositions rather than truth values. Insisting that semantic networks be allowed to represent only intensions suggests promising approaches to several knowledge representation problems which often lead to confusion. One of our goals is to devise a representation which is rich enough to maintain the subtle distinctions related to referential opacity and extensional equivalence. The purpose of the representation is to provide a substrate by which processes (e.g., programs, production rules, and inference rules) can operate.

The first problem we treat concerns indirect reference, originally raised by Frege (1892) and recently discussed by McCarthy (1979). McCarthy has put the problem into the following form:

> the meaning of the phrase *"Mike's telephone number"* in the sentence *"Pat knows Mike's telephone number"* is the concept of Mike's telephone number, whereas its meaning in the sentence *"Pat dialled Mike's telephone number"* is the number itself. Thus if we also have *"Mary's telephone number = Mike's telephone number,"* then *"Pat dialled Mary's telephone number"* follows, but *"Pat knows Mary's telephone number"* does not. (p. 129-130, italics in original)

Knowing is said to create an opaque context in its complement position and *dialling* is said to create a transparent context. The Uniqueness Principle suggests a solution strategy for this problem. We treat the concept of Mike's phone number and the concept of the number itself as distinct intensions, and thereby create a representional substrate with sufficient resolution to control inference processes which differentially apply in opaque and transparent contexts.

---

*The work presented here was partly supported by the National Science Foundation under Grant No. MCS78-02274 and MCS80-06314. We appreciate the comments and suggestions of the members of the Graduate Group in Cognitive Science and of the SNePS Research Group, both of SUNY at Buffalo. We also thank one of the anonymous reviewers for very extensive comments.

The second problem we treat is that of representing the concept of the truth value of a proposition. Since it is meaningful to talk about the truth value of a proposition independently of whether we know the proposition is true or false, we should have the ability to explicitly represent the truth value if we need to. For example, any semantic representation of the sentence "John knows whether he is taller than Bill" seems to explicitly reference the proposition underlying the sentence "John is taller than Bill." This seems clear after examining the paraphrase "John knows the truth value of the proposition John is taller than Bill."

Finally, we show how questions may be represented in propositional semantic networks. By this we mean how to represent the proposition that someone requested certain information of someone else. Such a proposition is contained in the sentence "I got mad because John asked me whether I was a boy or a girl."

HISTORY OF THE PROBLEMS

In this section, we briefly review the histories of three problems in knowledge representation for which we propose new solutions later in the paper. The main thrust of this paper is an investigation of the semantics of one kind of semantic network. That the theory leads to nice solution strategies for the three problems should be taken as a further explication of and support for the theory.

We can trace the emergence in artificial intelligence (AI) of the first problem to a paper by McCarthy and Hayes (1969) and it surfaced in its present form in later papers by McCarthy (1977, 1979). A well known philosophical solution to this problem, offered by Quine (1956) is to treat *knows* as an operator which creates an opaque context and to disallow substitution of equals by equals in an opaque context. Another approach has been adopted by Moore (1977, 1980). He encodes Hintikka's (1971) modal logic of knowledge within first-order logic and then builds an inference engine for that logic. An approach also using first-order logic is taken by McCarthy (1979) and Creary (1979). These researchers propose to view a concept as an object of discourse within the logic. They have one term which denotes Mike's phone number and another which denotes the concept of that phone number. This enables them to bypass the problem of replacement of equals by equals because the concept of Mike's phone number and the number itself are different entities. We differ from McCarthy in that we use one node to denote the concept of Mike's phone number and another node to denote the *concept* of the number itself.

One attempt to represent information about knowledge and belief in a semantic network was offered by Cohen (1978), but it is unclear how his treatment would relate to the issues raised in this paper. His primary goal was to show how a speaker's beliefs about a hearer's beliefs influences the speaker's planning of speech acts. Anderson (1977, 1978) has sketched a semantic network based procedure for processing referring expressions. Anderson's approach involves creating two distinct nodes in the network for the two concepts (i.e., the concept of Mike's phone number and its referent), and perhaps was inspired by Woods' (1975) important paper. He does not work out the details of representing belief, as they are not his primary goal. He does some interesting reaction time experiments which partly test the psychological reality of the Anderson-Woods proposal and these will be discussed later in the paper.

The solution we propose is in the spirit of Anderson but is more thoroughly articulated in the following ways: 1) we specify exactly how the nodes are related and which features of the representation trigger which kinds of inference processes; 2) we provide the solution with well articulated philosophical underpinnings; 3) we point out that these same philosophical assumptions provide an identical solution to some, at first glance, unrelated problems; 4) Anderson's experiments used stimulus materials which involved only transparent contexts. We suggest that the experimental results will not generalize to opaque contexts; and, 5) We point out that Anderson's model, if it is straight-forwardly extended, predicts some counter intuitive memory confusions.

The second problem, of representing the notion of truth value, has not received much attention in the AI literature, probably because the problem domains which have been attacked by AI researchers have allowed domain specific solutions. We feel the problem deserves more attention for two reasons. First, the notion of truth value in general and the notion of the truth value of a specific proposition in particular can be objects of thought themselves. Since propositional semantic networks are supposed to be knowledge representations in which every concept (object of thought) is represented by a node, the notion of truth value should also be so represented. Second, having decided that truth values of propositions might be represented by nodes in a semantic network, we quickly find where such nodes can be useful. For example, utterance (1)

(1) John knows whether he is taller than Bill.

can be taken as an assertion that mentions the truth value of the proposition that John is taller than Bill without taking a position on whether it is true or false. Thus, in order to represent (1) as a proposition about a truth value, we need to be able to represent a truth value independently of the specific value (true or false) it happens to be. An alternative solution used by Allen (1979) and Cohen and Perrault (1979) involves specifying the meaning of "knowing whether" as a disjunction of correct beliefs. For instance, the disjunction (2) could serve as the representation for "John knows whether *P*."

(2) (*P* & John believes *P*) or
(not *P* & John believes not *P*)

We claim that this representation uses the intensions *P*, John believes *P*, not *P*, and John believes not *P*, but it does not use the intension of the truth value of *P*, which is, we believe, the intension of "whether *P*."

A simpler example which we shall use throughout this paper is shown in (3) with its disjunctive reading in (4).

(3) John holds an opinion about whether he is taller than Bill.
(4) John believes he is taller than Bill or John believes he is not taller than Bill.

This disjunctive solution does not generalize to some embedding verbs such as "wonder," which is shown as sentence (5)

(5) John wondered whether *P*.

Perhaps (5) could be paraphrased as (6a) and thence, via the disjunctive reading of "knows whether" as (6b). We feel, however, that a better paraphrase is (6c), which mentions the concept of a truth value explicitly.

(6a) John wanted to know whether *P*.
(6b) John wanted to know *P* or to know not *P*.
(6c) John was curious about the truth value of *P*.

Our only claims are that a node for the concept of a truth value should explicitly exist in the network and is useful for certain representation tasks. We agree that (3) and (4) are logically equivalent and this can be captured by inference rules of some kind. However, logical equivalence is not the same as intensional identity.

The last of the three problems we discuss, that of representing questions, acquires salience because of our position that propositional semantic networks represent only intensions and the combined facts that: 1) intensions are meant to correspond to individual concepts or propositions, and 2) questions are neither individual concepts nor propositions. Other representational schemes can trivially represent questions by tagging the symbol structure which describes the content of the question as being a question and not an assertion. For instance, Schank (1975) represents yes-no questions in conceptual dependency diagrams by indicating that the mode of the diagram is a question, rather than an assertion. Wh-questions are indicated by placing the symbol "*?*" in the question slot. That is, the question "Where is the salt?" would be represented as something like (LOC SALT *?*). A propositional network would interpret the notation (LOC SALT *?*) as a proposition stating that the salt is at the question mark, or a proposition involving a free variable.

It happily turns out that by viewing yes-no questions as enquiries about truth values we can immediately represent them. Later in the paper, we attempt to generalize this solution to wh-questions as well.

THE THEORETICAL FRAMEWORK

What Does a Semantic Network Model?

The first issue we must be clear about is what we intend a semantic network to model. We first exclude two possibilities.

One possibility is the real world. This would require nodes to represent objects in the world (as opposed to individual concepts) and facts about such objects (as opposed to propositions). Although some people might be interested in semantic networks as models of the world, we are not.

A second possibility is that a semantic network models a corpus of natural language text, or perhaps that a semantic network is a data structure in which a text is stored and from which pieces of the text can be retrieved easily. In this case, nodes of the network would represent words, lexemes, morphemes, strings, phrases, clauses, sentences, and so forth. In Woods (1975), for example, it is argued that the semantic network representation of "The dog that had rabies bit the man" must distinguish between the proposition of the main clause, "The dog bit the man" and the proposition of the subordinate clause, "The dog had rabies." Woods proposed this for semantic reasons (p. 62). Had he made the proposal for purely syntactic reasons, then he would have been representing sentences as opposed to meanings. Although we feel that semantic networks can be used to model natural language text (cf. Shapiro & Neal, 1982), this is not the use with which we are concerned in this paper.

A third possibility, and the one that we are concerned with, is that a semantic network models the belief structure of a thinking, reasoning, language using being (e.g., a human). In this case, nodes represent the concepts and beliefs such a being would have. The point is that these concepts are intensions rather than extensions.

This is not, perhaps, the majority view of researchers in "knowledge representation." In their survey of knowledge representation research, Brachman and Smith (1980) asked researchers, "between what two things do you envisage the 'representation' relationship?" (p. 68). They report, "The one interesting thing to be said in summary, it seems, is that the phrase which we use as a commonplace label of our endeavor—'the representation of knowledge'—is perhaps surprisingly not taken in a particularly literal sense. . . what was considered to be 'represented', typically, were various kind of object [sic] in the world" (p. 71). Nevertheless, in answer to the question, "Would you characterize your inquiry as primarily an investigation into the structure of *language*, or more as an investigation into the structure of thought? . . . *The great majority gave their prime allegiance to the study of thought*" (p. 71).

Intensions

The term "intension" derives from Frege's (1892) term "sense". He was concerned with the relationship between *equality* and the meanings of designating expressions in a language. The fact that (7a) is true and (7b) is false illustrates the problem which concerned Frege.

(7a) Necessarily, the Morning Star is the Morning Star.
(7b) Necessarily, the Morning Star is the Evening Star.

Frege took this as evidence that the designating phrases "the Morning Star" and "the Evening Star" do not have the same meaning, even though they denote the same object. If the expressions were equal then (7a) and (7b) should have identical meaning. Frege used the term "sense" of an expression to intuitively correspond to the meaning of that expression, and he used the term "reference" to correspond to the denotation of the expression.

Carnap (1947) attempted to formalize Frege's notion of the sense of an expression as a function from possible states of affairs to denotations. The function was called an "intension" and the denotation was called an "extension" (cf. Dowty, Wall, & Peters, 1981, p. 145). The approach was refined through Kripke's (1963) semantics of necessity and Montague's intensional logic (p. 145).

When we say that nodes of a semantic network represent intensions as opposed to extensions, we mean sense as opposed to reference. Additionally for the purposes of this paper, we will not view intensions as functions, although it might be helpful to do so in the future. When we say that nodes of a semantic network represent intensions we mean intension as Frege (1892), McCarthy (1979), and Creary (1979) view intension, as opposed to Carnap (1947), Kripke (1963), Montague (cf. Dowty, et al.), or Moore (1980). We take intensions to correspond to concepts, ideas, objects of thought, or things which can be conceived of.

The Need for Intensional Representations

Woods (1975) appears to have been the first to emphasize that some nodes of a semantic network should represent intensions. One reason for this is to enable the cognitive agent being modeled to conceive of things which do not exist in the actual world such as unicorns and Santa Claus. Although unicorns do not exist, they can be reasoned about. Indeed, the reader can say, as McCawley (1980) points out, how many horns a two-headed unicorn would have. Thus, we should be able to describe within a semantic network any conceivable concept, independently of whether it is realized in the actual world, and we should also be able to describe whether in fact it is realized.

Returning to the Morning Star—Evening Star example, Woods (1975) concluded that, "there must be two mental entities (concepts, nodes, or whatever) corresponding to the two different intensions, morning star and evening star. There is then an assertion about these two intensional entities that they denote one and the same external object (extension)" (p. 50). Woods continues with the observation that "there must be some nodes in the network which correspond to descriptions of entities rather than entities themselves. . . . We have to decide how to tell the two kinds of nodes apart" (p. 68). Semantic network theorists have not universally agreed with this position. Schubert, Goebel, and Cercone (1979), for instance say, "We take the position that terms (nodes, subnets) *already have both extensions and intensions*" (p. 128, italics in original). Brachman (1977) takes a position on the other side of Woods, stating, "*Semantic networks are representations of the intensions of natural language designators*" (p. 139, italics in original). Yet Brachman still allows some extensional information in his networks: "some of the operations in the network scheme are purely intensional, while others are not" (p. 150).

We want to go even further than Brachman and say that all information in the network is intensional. The only reference to extensions will be propositions stating that two intensions pick out the same extension in some world, such as the proposition that the Morning Star is the Evening Star in the actual world.

The Absence of a Need for Extensional Representation

If, as Woods pointed out, at least some nodes represent intensions, do any represent extensions? Indeed, should every node be seen as having both an intension and an extension as Schubert et al. claim and as most philosophers usually treat designating expressions? Our answer derives from what we take our networks to model (see above). If a network modeled the real world, then a node would represent (or denote) an extension, but we take our networks to model conceptual belief structures.

A node that represents only an intension carries no commitment that an object realizing the intension exists in the real world. The standard translation of "The present king of France is bald" into a logical notation seems to require asserting the existence of the present king of France.

(EXISTS x) (x is the-present-king-of-France & x is bald).

However, this is because of the normal extensional interpretation of statements in standard logic. A constant node in a semantic network is like a Skolem constant derived from a extensionally quantified variable that asserts only the existence of the intension.

STRUCTURAL IMPLICATIONS OF INTENSIONAL REPRESENTATION

The Need for Co-referential Propositions

If a semantic network has a node for the intension of the Morning Star and a different node for the intension of the Evening Star, what should be done when the assertion is made that the Morning Star is the Evening Star? If the nodes were merged by transferring all of the arcs from one node to the other and eliminating the first, then this would eliminate the distinction between the two concepts and make it impossible to represent the sentence "John did not know that the Morning Star is the Evening Star" differently from the sentence "John did not know that the Morning Star is the Morning Star." The solution Woods (1975) proposed is to add a node to the network representing the proposition underlying the sentence "The Morning Star is co-referential with the Evening Star in the actual world." The co-referential proposition can be used by the system's reasoning processes to infer that certain beliefs about the intension of the Morning Star can be transferred to (and from) the intension of the Evening Star. This will be discussed further below.

Order Dependency

The set of nodes existing in a network will depend not only on what information is presented to the network but also on the order of presentation. We shall call this property *order dependency.* Consider Russell's (1906) example, "George IV wished to know whether Scott was the author of *Waverly*" (p. 108). This apparently came about because, even though Scott was well known at the time, *Waverly* was published anonymously. George the IV had an intension for Scott and an intension for the author of *Waverly* but did not know whether they were extensionally equivalent. A semantic network simulating George IV, or just recording this fact, would need a different node for each of these intensions. But what does this imply for how we should represent the information that Scott wrote *Ivanhoe* (assuming that the sentence "Scott wrote *Ivanhoe*" is our first introduction to the novel)? Must we represent this as two propositions, one for asserting the co-referentiality of the intension for Scott and the intension for the author of *Ivanhoe,* and another for asserting that the author of *Ivanhoe* wrote *Ivanhoe.* Our answer is that the cognitive system creates intensions only as needed for storing information about them. A separate concept was needed for the intension of the author of *Waverly* only because there was thought of that author before an identification was made with any previous intension. Psychological evidence for this analysis has been provided by the work of Anderson and Hastie (1974), Anderson (1977, 1978), and McNabb (1977). This will be discussed in a later section.

This theory would predict that a sentence such as "George IV thought that the author of *Waverly* was older than Scott," which requires two intensions having the same extension, would be harder to understand by someone whose first introduction to the intension of the author of *Waverly* was via the statement, "Scott wrote *Waverly*" than by someone who had already thought about Scott and the author of *Waverly* independently. The second person would already have two intensions to use. The first person would at first access the same intension for both Scott and the author of *Waverly,* but would then create at least one new concept for the intension of the author of *Waverly.* We call this process *splitting* and will return to a more detailed description of it later. This phenomenon would explain part of the cuteness of the example "Shakespeare's plays weren't written by him, but by someone else of the same name" (Hofstadter, Clossman, & Meredith, 1980).

Opacity as the Norm

A system that conforms to the Uniqueness Principle does not need the substitutivity of equals for equals as a basic reasoning rule, because no two distinct nodes are equal. Co-referentiality between two nodes must be asserted by a proposition. It requires inference rules to propogate assertions from one node to another node which is co-referential with it. Thus, intensional representation implies that referential opacity is the norm and transparency must be explicity sanctioned by an inference process, unless nodes are "merged." Merging will be discussed in a later section.

Connections with Reality

The main objection to exclusive intensional representation seems to be that if nodes represent only intensions, how could any alleged understanding system so based have any connections with the outside world? To consider this question, we must endow our modeled cognitive agent with sense and effector organs. We will look at a robot system with sight and manipulators.

The robot needs a perceptual system in which some node, set of features, and so forth is triggered consistently when a given object is seen. If it is to communicate reasonably with people about perceptual topics, it must be able to make approximately the same perceptual distinctions that we make. These perceptual nodes need not extensionally represent the objects that trigger them. *The perceptual nodes can be connected to the semantic-conceptual nodes. This allows the robot to "recognize" objects, although it could be fooled.*

The robot also needs effector organs that operate the manipulators consistently, and connections between some semantic-conceptual nodes and the effector organs so that it can operate its manipulators in a manner dictated by its reasoning (decide what to do).

Sensors and effectors, by supplying a connection between some node in the semantic network and some object(s) in the actual world, finally provide referents for the node in one particular world. However, these referents are only exemplars of the concept represented by the node. The significance of the node remains its intension.

The Meaning of the Nodes

There are a number of formalisms compatible with the definition of propositional semantic network presented at the beginning of this paper. The formalism used here contains only propositions and individual concepts. We will not specify a formal semantics for the network structures because the meaning units of the network, the nodes, violate Frege's Principle of Compositionality (cf. Dowty, Wall, & Peters, 1981; p. 42). If a formal language obeys the principle of compositionality, then for any non-atomic expression in the language, there is a set of rules which enable one to determine the expression's meaning from the meaning of the expression's subexpressions. In turn, the meaning of the subexpressions can be determined on the basis of the meaning of the sub-subexpressions, and so on, recursively.

There are two properties of a formal semantic system that must hold before one can even think of writing the above mentioned rules. The networks which appear in this paper have neither of those properties. First, in order for the above mentioned recursion to terminate, the language must have atomic expressions, or semantic primitives. However, the only primitives used in the networks are arc labels, which are non-conceputal, serve a syntactic function, and have no meaning. Second, in order for the principle of compositionality to apply at all, the meaning of an expression must not change when it is embedded in another expression. The meaningful expressions in our network, the nodes, get their meaning from the expressions they are embedded in (in addition to getting meaning from their subexpressions) and thereby change their meaning whenever they are embedded in a new expression.

Although our formalism does not obey Frege's principle, its characteristics which violate the principle can be found in Quillian's (1968) writings. Quillian, as Brachman (1979) has pointed out, is most known for suggesting that semantic information be stored in a subclass-superclass hierarchy so that properties could be inherited from superclass to subclass. There is, however, a less well known aspect of his writing that is relevant to our discussion. The next two quotations, taken from Quillian (1968), seem to argue for a data structure that violates the principle of compositionality for the two reasons stated above. In the quotations, we will interpret the term "word concept" to mean node. First, Quillian argued that there are no primitives:

> there are no word concepts as such that are "primitive." Everything is simply defined in terms of some ordered configuration of other things in memory. (p. 239)

Second, Quillian argues that the meaning of a node is determined by a search process which begins with that node:

> *a word's full concept is defined in the memory model to be all the nodes that can be reached by an exhaustive tracing process, originating at its initial, patriarchal type node*... (p. 238, italics in original)

This implies that, as additional structure is added to the network, it changes the modeled cognitive agent's understanding of the concepts represented by every node connected to the added structure. Martin (1981) has applied the term *decompositional* to nodes whose meaning is dependent upon what they are constituents of. Because of this extreme form of decompositionality, the network described here will be inclined to memory confusions precisely because of these decompositional meanings. The mechanics of these confusions will be described in a later section. Anderson (1978, p. 51, Figure 1) appears to use decompositional meanings. It is an empirical question as to whether humans are inclined to make the above mentioned confusions. It may be the case that we will have to augment the network notation so that the meanings are more stable and conform to the second prerequisite for compositionality. Woods' (1975) arguments for the need to distinguish between the definitional information associated with a node and the assertional information associated with a node are relevant here.

THE REPRESENTATION CONSTRUCTS

Here we introduce the network notation used in this paper. It is not the only network notation compatible with semantic networks but uses perhaps the fewest constructs. The network will have two kinds of nodes which we will call base nodes and proposition nodes. The proposition nodes represent propositions and the base nodes represent individual concepts. All arcs are directional and are either ascending or descending, and come in pairs such that one arc of a pair is descending and the other is ascending. In the diagrams of this paper, only the descending arc of a pair is shown but there is always an ascending arc (corresponding to the inverse of the descending arc) which is not shown. A node is a base node if it has no descending arcs emanating from it other than a LEX arc. Otherwise, a node is a proposition node.

The LEX Arc

The purpose of the LEX arc is to form an access route from an individual concept to a word in a lexical memory. The LEX arc is a device used in this paper only to simplify the presentation. One drawback of the LEX arc is that its function is different than the function of the other arcs. This is because the object which it points to is not an intension, but rather an entity in lexical memory. Using the LEX arc to associate a name, say "John", with a concept, say the intension of John, does not give the network knowledge (in a conceptual sense) that John's name is "John," and we cannot tell the network that it has the wrong name for the person it thinks is named "John." A more consistent way to do this would be to have an intension for the word "John" as well as for the person John, and then use a proposition to assert that the intension for "John" is the word for the intension John. Any links to lexical memory would then emanate from the intension for "John" and not the intension for John. For the purposes of exposition we are representing an approximate distinction between a concept and a word denoting the concept but we would not be able to describe the process of acquiring word meanings with this scheme. So the current scheme is being used only to make the diagrams and presentation simplier.

The important things to know about the LEX arc are that: 1) it points from a concept to a word, and 2) the node it emanates from is an individual concept and not a proposition.

Other Descending Arcs

Other than the LEX arc, descending arcs always point to arguments of propositions and the nodes from which they descend represent propositions. A proposition node which has no descending arcs coming into it is said to be a non-dominated node and represents a proposition that is believed by the system. It is also necessary to be able to attach erasable, nonconceptual, assertion tags to proposition nodes which are dominated. The purpose of these tags is to indicate that the system believes those propositions. If a dominated proposition does not have a belief tag then the system has no opinion about its truth. The belief set of the semantic network consists exactly of the nodes which are non-dominated or which are tagged as true by an assertion tag.

Extensional Equivalence

There is a case frame, called the equivalence case frame, which is used to represent extensional equivalence. This case frame is used to assert that two distinct intensions have the same extension in the real world. It does not say what the extension is but rather only that they have the same extension. Thus "The Morning Star is the Evening Star" is represented in the simplified network structure of Figure 1 and can be paraphrased as "The conceptual object denoted by 'Morning Star' is extensionally equivalent to the conceptual object denoted by 'Evening Star'." The reason the case frame has a special status is that it will interact with propositions involving knowing and believing in ways that other propositions will not interact.

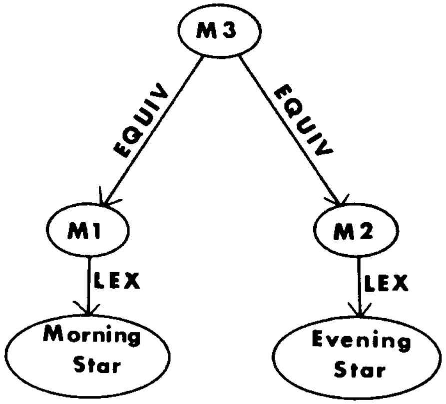

Figure 1. A representation (node M3) for the propositional information in the sentence "The Morning Star is the Evening Star."

Knowing and Believing

We will also use two different *know* relations. One will be used for knowing a fact as in the sentence "John knows that he is taller than Bill." The other will indicate familiarity with a concrete or abstract object as in knowing a person or game, or as in the sentence "Pat knows Mike's phone number." We will call the sense of knowing a fact "know 1" and the sense of being acquainted with a thing "know 2". We justify these relations on the basis of the introspection that knowing an object is different than knowing a fact. Similarly, we will use two different *believe* relations, called "believe 1" and "believe 2," to correspond to the two different *know* relations.

As stated earlier a propositional semantic network models the belief system of a cognitive agent, which we will call "the system." Whenever any of the relations involving knowing or believing appear in the network, they will represent beliefs about beliefs.

We shall treat know1 only as correct belief rather than justified, correct belief for the following reason. *If a cognitive agent believes that someone* knows *a facts, then the agent believes: 1) that the someone believes the fact; 2) the fact is true; and 3) that the someone believes the fact for the* right reasons (as far as the cognitive agent can tell). The third stipulation about "right reasons" is necessary to rule out belief for the wrong reasons such as superstition or guessing. For example, there is at least one person who believes that white can force a win in chess, and there is at least one person who believes that white cannot. Therefore we know there is a person who has a correct opinion about whether white can force a win. Nonetheless, no one currently knows whether white can force a win. Since we are not able to specify what "believing for the right reasons" means, we will discuss knowing only in the sense of correct belief.

McCARTHY'S TELEPHONE NUMBER PROBLEM

The Main Example

We now address ourselves to McCarthy's telephone number problem. What follows is our representation, and its rationale, for sentence sequence (8).

> (8) Pat knows Mike's phone number; Ed dials Mike's phone number; and, Tony correctly believes that Mike's phone number is the same as Mary's phone number.

This sentence sequence illustrates the distinctions upon which McCarthy's example focuses. On the basis of (8) the system should conclude that Ed dials Mary's phone number, yet it should not conclude that Pat knows Mary's phone number. Furthermore, it should not conclude that Tony knows either Mike's or Mary's phone number; however, if the system is subsequently told that Tony knows one of the numbers in particular, it should conclude that Tony also knows the other number. Figure 2 shows the network representation of the information contained in (8). (Note: Although it does not appear in the figure, assume M5 has been tagged as true by an assertion tag.)

Critical to our discussion is an explication of the system's understanding the concept of node M7 of that figure. Node M7 is:

> Something which Pat knows, and
> something which Ed dials, and
> something which is Mike's phone number, and
> something that is co-referential with Mary's phone number, and
> something that Tony knows is co-referential with Mary's phone number.

In order to explicate the node's full concept, it was necessary to traverse the entire network, and the network is the sum total of the system's beliefs.

The assertion that node M7 is co-referent with Mary's phone number is just as much a part its meaning as the assertion that it is Mike's phone number. However, we do not want the system to decide that Pat knows Mary's phone number, so how do we avoid this? This is where extensional equivalence links acquire their special status. If the system believes that

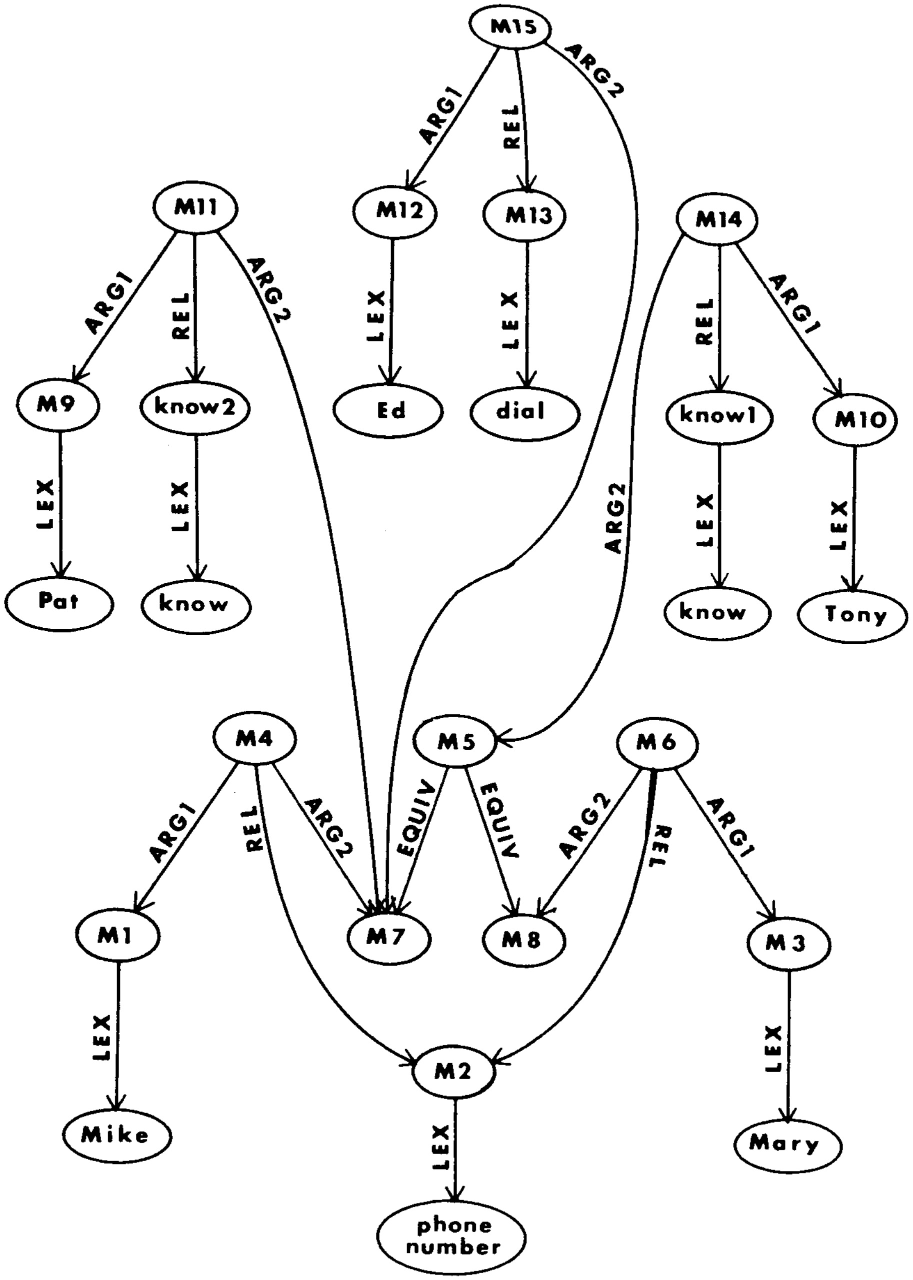

Figure 2. A representation for the information in the sentence sequence: "Pat knows Mike's phone number; Ed dials Mike's phone number; and, Tony correctly believes that Mike's phone number is the same as Mary's phone number."

some agent knows some intensional concept, then it will assume that the agent knows all of the propositions asserted about the intensional concept except those propositions which consist of the extensional equivalence case frame (e.g., M5) or which contain the extensional equivalence case frame (e.g., M14). So the system will assume that Pat knows M4, M11, and M15. Now the system could be mistaken; Pat need not necessarily know M15. This performance characteristic would predict corresponding kinds of memory confusions in humans. Further discussion of this appears later in the paper in the section titled "Merging and Splitting Nodes."

Further Examples

To further illustrate how we would use the representation, we will look at two more sentences, and then return to the question of transparency verses opacity. Consider (9) below.

(9) Mike has two phones. One is 831-1234. The other is 831-4321.

Figure 3 shows how we represent it. The next example illustrates how we handle designators which fail to refer, as seen in (10) below.

(10) Mike doesn't have a phone.

There are two ways to represent this utterance. The first involves the intension of nothing (for a discussion of this intensional object see, Heath, 1967). It is the notion of non-existence. The representation is depicted in Figure 4, and if one of the non-dominated nodes were submitted to a language generator (node M4 or M5), it might produce the sentence "Mike's phone number is non-existent." The other way employs universal quantification and, as shown in Figure 5, asserts that for all x, x is not Mike's phone number. The quantification arc is taken from the SNePS semantic network formalism (Shapiro, 1979a).

Transparent Relations Propogate by Inference

We now return to describing how the representation for McCarthy's telephone number problem can support processes which simulate referential transparency and opacity. Opacity is the norm because for instance, in the situation of sentence sequence (8), Mike's phone number and Mary's phone number are distinct intensions which are extensionally equivalent; but extensional equivalence is not equality, so we do not encounter the problem of substitution of equals for equals. Transparency, however, requires some sort of inference process. What is needed is an inference or production rule which propogates assertions across equivalence arcs, but which has an ante-

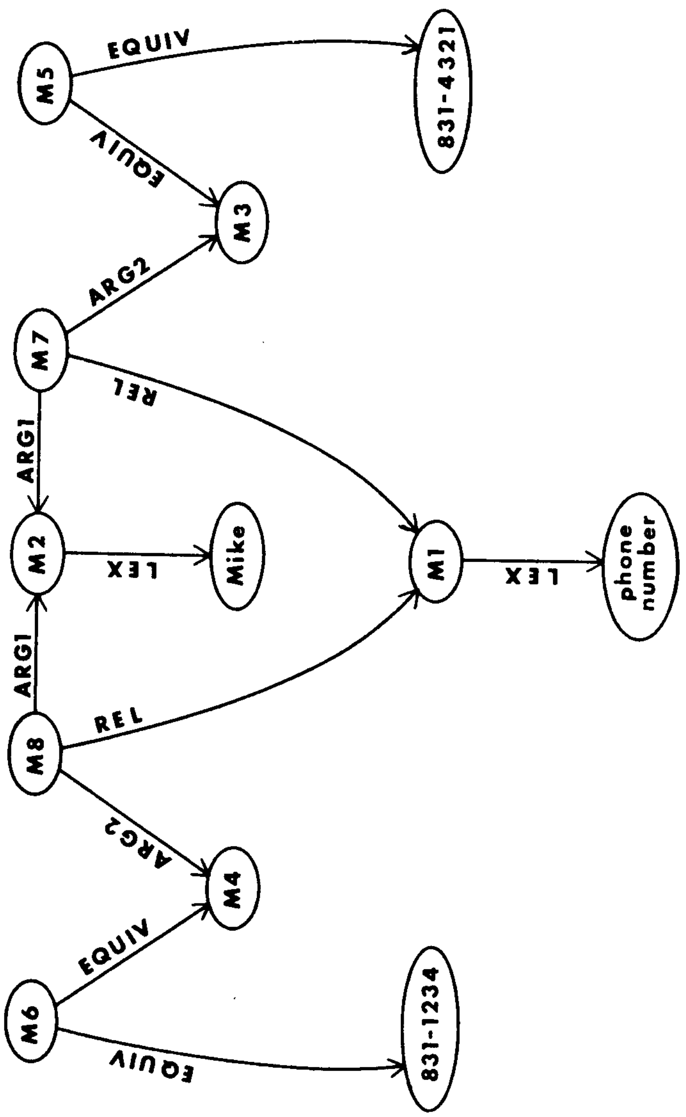

Figure 3. The representation for Mike having two phones.

cedent trigger pattern that matches only assertions involving transparent operators. This entails that the system have the conceptual knowledge that the relation *dial* is transparent but *know* is not. Such an inference rule would then enable the system to add to its data base that Ed dials Mary's phone number from the information given in (8); yet the system would not add that Pat knows Mary's phone number, because *know* is not transparent and would not satisfy the trigger pattern of the inference rule.

There are situations in which assertions involving opaque operators can also propogate across equivalence paths. Tony in (8) knows that Mike's phone number and Mary's phone number are extensionally equivalent, so any assertion with Tony as the agent and Mike's phone number (node M7) as the object, regardless of whether it involves an opaque operator, should be able to propogate across that equivalence path. In short then, if a cognitive agent knows that two concepts are extensionally equivalent, then for that agent all operators are transparent with respect to that path. The appendix contains examples of inference rules which propogate assertions across an equivalence path. No claim is made that they are complete.

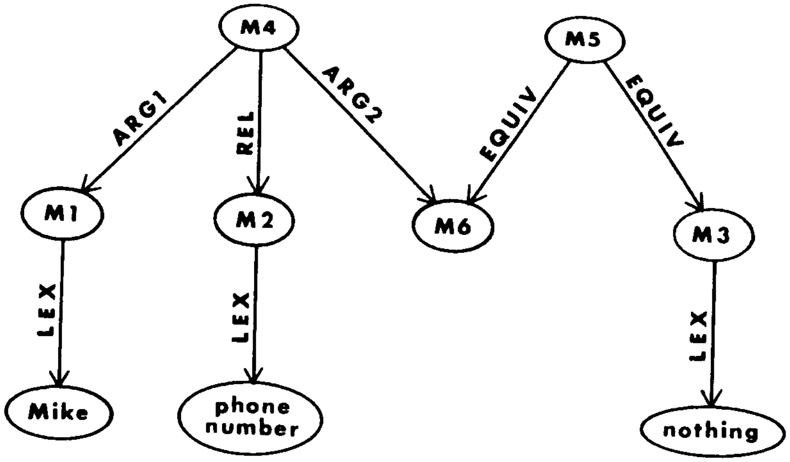

Figure 4. The representation for Mike not having a phone.

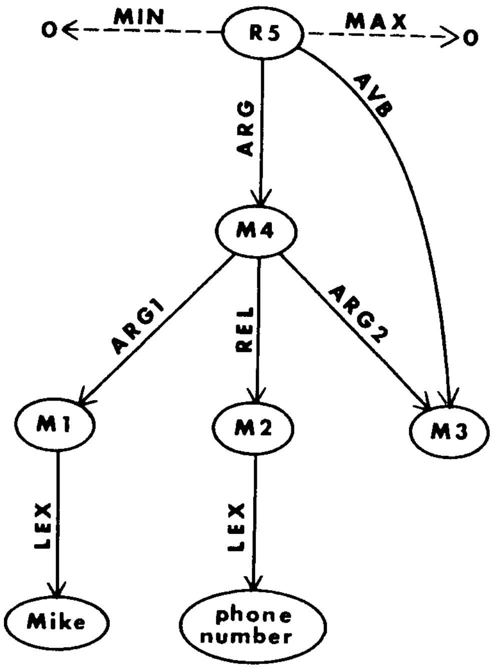

Figure 5. Another representation for Mike not having a phone.

REPRESENTING TRUTH VALUE

This section makes several points. First, the concept of the truth value of some proposition is intensionally distinct from its particular truth value (true or false) and at times it is necessary to make this distinction explicit in the representation. Second, the same information can be represented by two different configurations of concepts. This illustrates what we mean by order dependency. And third, we contrast the behavior of the system's beliefs about believing with its beliefs about knowing (correct belief).

Figure 6 shows the representation for sentence (11), with the proviso that the system has already wondered about the truth value of the proposition underlying the sentence "John is taller than Bill." An implication of our discussion of order dependency is that a distinct intension for the truth value of a proposition will be created only if the system wonders about the truth value of that proposition before actually learning its truth value.

(11) It is true that John is taller than Bill.

This treatment embodies the claim that the extension of a proposition is its truth value. Node M8 is the intension for the truth value of the proposition underlying the sentence "John is taller than Bill." Node M6 represents the individual concept of truth. When John wonders whether he is taller than Bill as in sentence (5), he is trying to determine the co-reference of node M8.

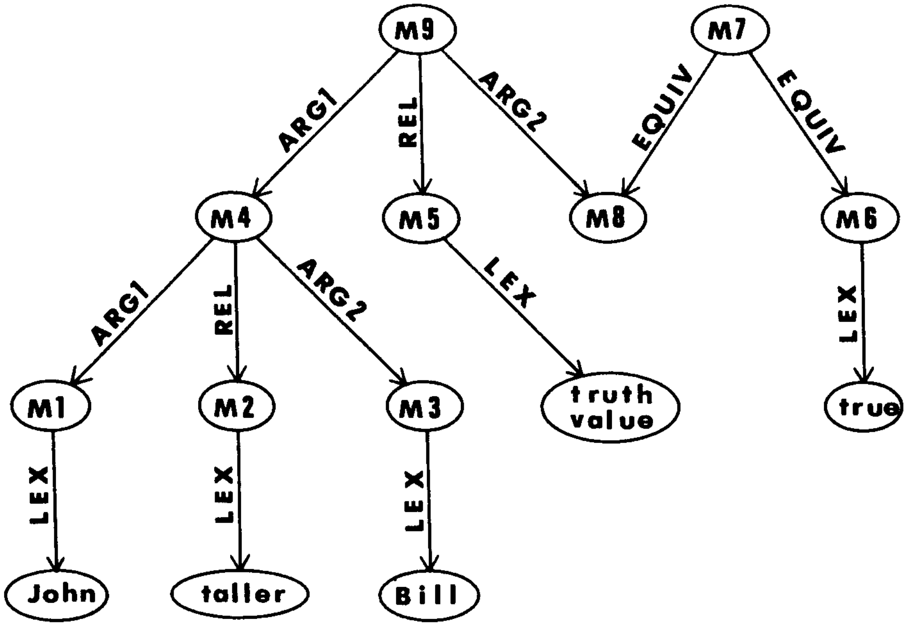

Figure 6. The representation for the propositional information contained in the sentence "It is true that John is taller than Bill."

If Tony believes, as in sentence (1), John knows whether John is taller than Bill, then Tony believes John's description of node M8 is complete even though Tony's is not necessarily complete.

We emphasize that the configuration of nodes used to represent (11) depends on the order in which the system thinks of them. If someone is directly told that John is taller than Bill and does not have to wonder about it, then the representation would be much simpler, as shown in Figure 7, in which node M4 represents the proposition underlying the sentence "John is taller than Bill."

Since (11) is an extraposed version of the sentence "That John is taller than Bill is true," we suspect it carries the presupposition that the listener has in fact wondered about the truth value of the embedded proposition, and so the representation depicted in Figure 6 is the correct one for this sentence. The EQUIV-EQUIV case frame indicates that two intensions already represented in the network are extensionally equivalent in the actual world. Node M6 of Figure 6 would presumably already exist in anyone's memory who has some notion of truth. Mode M8 would exist in people's memory who have wondered about the truth value of "John is taller than Bill" before learning its actual truth value.

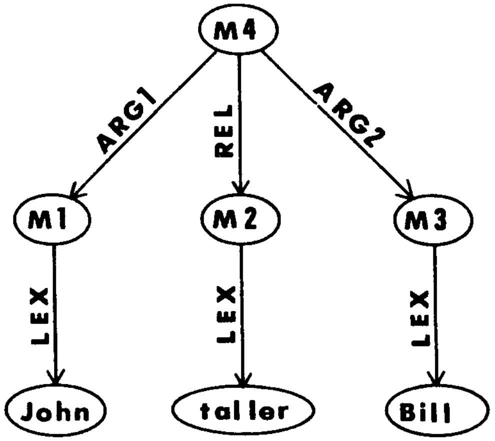

Figure 7. The representation (node M4) for the propositional information contained in the sentence "John is taller than Bill."

The reader might entertain the possibility of using the above technique as a general method of asserting propositions in the network. This is not possible because it would lead to an infinite regress of assertion embeddings. If the EQUIV-EQUIV case frame were necessary to assert a proposition then Figure 6 would not assert utterance (11) because node M7 has not been asserted by the EQUIV-EQUIV case frame. The convention of taking top-level nodes as being asserted alleviates this problem.

In order to represent sentence (3), duplicated below, we use "believe 2." As already mentioned, it means to be familiar

(3) John holds an opinion about whether he is taller than Bill.
(4) John believes he is taller than Bill or John believes he is not taller than Bill.

with, or apprehend, some intension. Figures 8 and 9 contrast our representations of sentences (3) and (4). In Figure 9, sentence (4) is represented by node M8 which is a proposition involving exclusive-or. A feature inherent in the use of "know 2" is that the task of appraising in exactly what manner the agent of the "know 2" relation is familiar with, or apprehends, an individual concept requires the use of inference rules. "Know 2" says nothing in itself except that the agent is familiar in some domain specific way with the individual concept that is the object of the relation. *There must also be a*

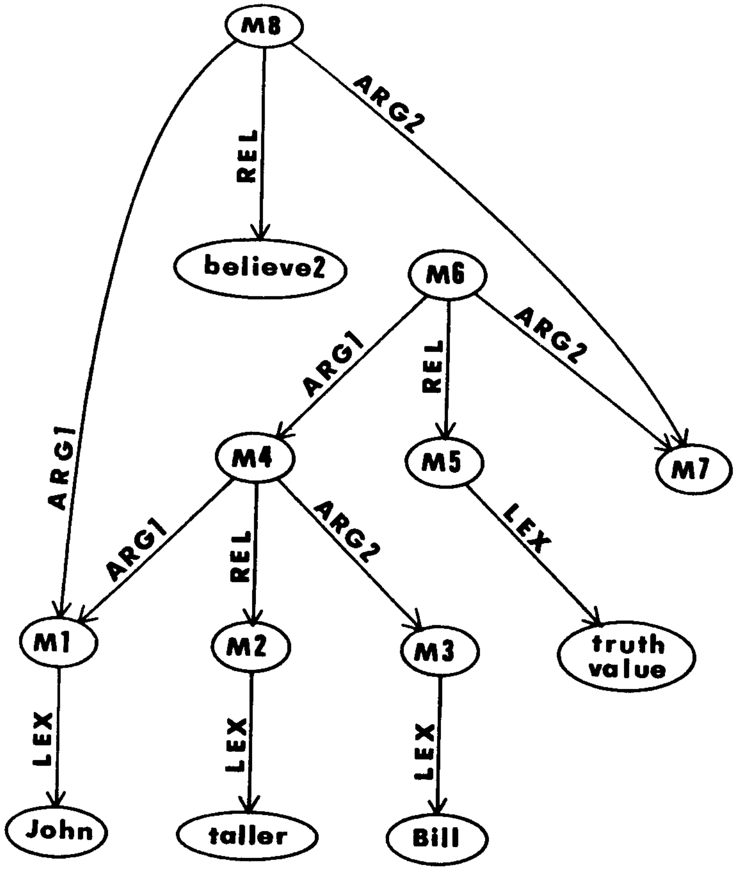

Figure 8. The representation (node M8) for the propositional information contained in the sentence "John holds an opinion about whether he is taller than Bill."

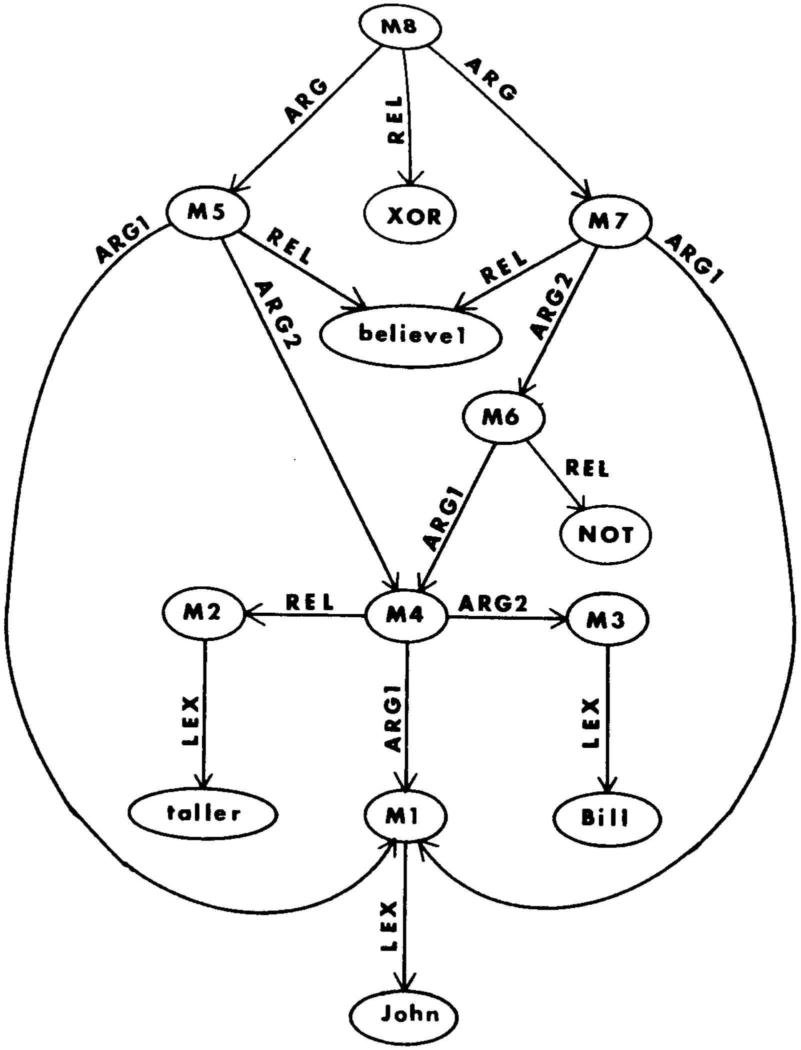

Figure 9. The representation (node M8) for the propositional information contained in the sentence "John believes he is taller than Bill or John believes he is not taller than Bill."

way for the specific facts that an agent knows about the concept to be independently asserted. An illustration of this point is depicted in Figures 10 and 11. Node M8 in Figure 10 does not directly represent

(12) John knows that he is taller than Bill.

sentence (12), but rather entails sentence (12). Node M5 in Figure 11 does directly represent (12). Believing that *P* and having a correct opinion about whether *P* are two different intensions and therefore should be represented by two distinct nodes. The network structure of Figure 10 does not explicitly assert what John's opinion actually is. Using the inference rule (13) below, the information in Figure 11 could be deduced from the information in Figure 10.

(13) If the system believes1 that some cognitive agent knows2 whether some proposition *P* and the system believes1 that *P*, then the system can conclude that the agent believes1 that *P*.

The following example illustrates the behavior of this kind of system. Suppose you were a college freshman majoring in computer science, and suppose you believed (incorrectly) that pi was a rational number. Naturally, you would also believe that your computer science professor also knew whether pi was rational. By applying the above inference rule, you would erroneously, but properly, conclude that your professor believed that pi was rational.

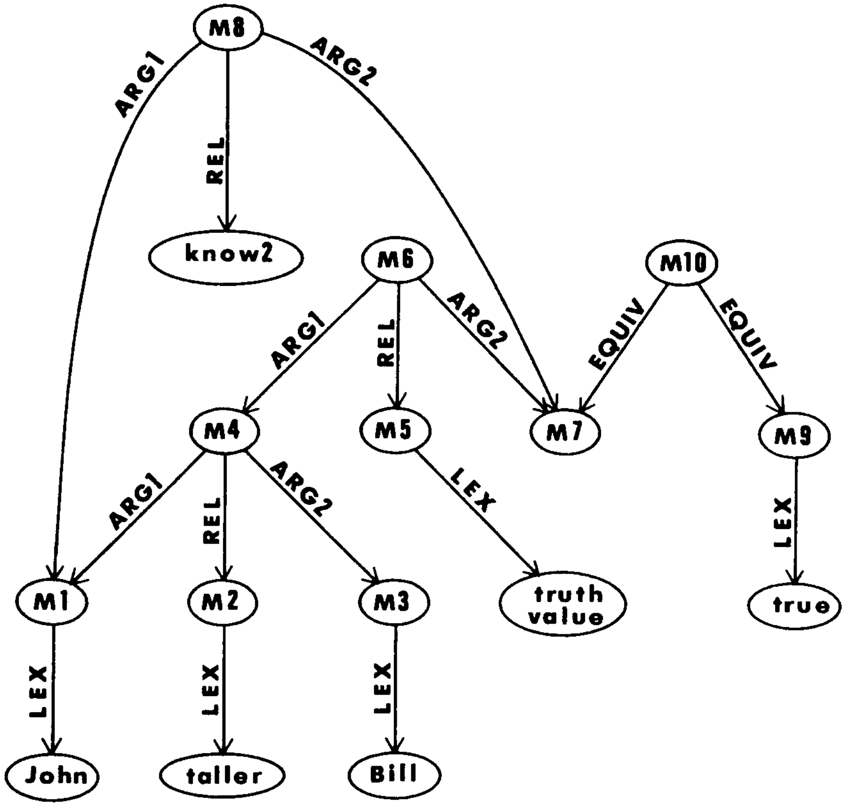

Figure 10. The representation for the sentence sequence: "John holds a correct opinion about whether he is taller than Bill; it is true that John is taller than Bill."

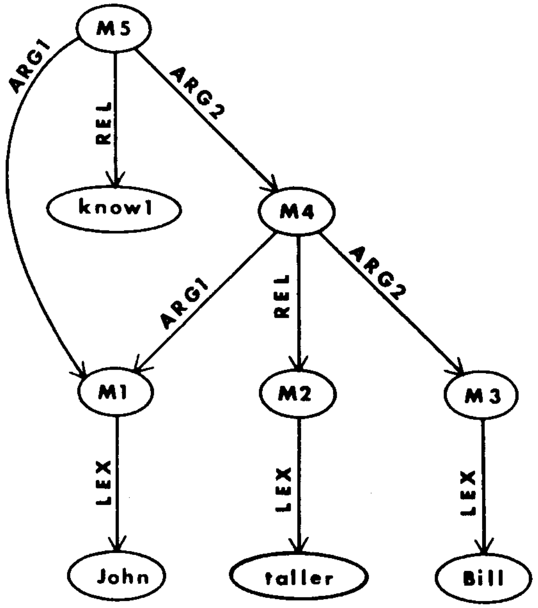

Figure 11. The representation (node M5) for the propositional information contained in the sentence "John holds a correct opinion that he is taller than Bill."

REPRESENTING YES-NO QUESTIONS

In order to support the entire range of discourse processing it is necessary for a knowledge representation in general, and a semantic network in particular, to have the ability to represent questions. Although it is possible for a network which does not represent questions to interface with a processor that enables it to answer questions (e.g., an ATN parser-generator; Shapiro, 1979b), the network itself would intrinsically not have the ability to remember which question were asked. The need to represent questions is illustrated by sentences of the sort "Since John asked me whether he was taller than Bill, I assumed he had never met Bill." In order to represent this sentence, it is first necessary to represent the question embedded within the subordinate clause.

We now turn to the problem of representing a yes-no question in a propositional semantic network. A question is not an individual concept, nor is it a proposition, but it must be represented in a propositional network by some configuration of individual concepts and propositions. *It is not the case that a question can be paraphrased into a propositional form. Consider the line of dialog below (14) and the attempt to propositionalize it (15).*

(14) Mary said, "John are you taller than Bill?"
(15) Mary asked John whether he was taller than Bill.

The paraphrased form superficially resembles a proposition except for a constituent centered around "whether," which cannot be propositionalized. There is a way to circumvent the obstacle by noticing that the sentence resembles sentence (5). We can generate (16) from (15) by replacing "ask" with "wonder" as shown below.

(16) Mary wondered whether John was taller than Bill.

Our treatment of "wonder" described earlier applies here and can be extended to "ask." Sentence (15) is an enquiry about the truth value of the proposition underlying the sentence "John is taller than Bill." The literal interpretation of all yes-no questions can be handled in this manner. Our representation of (15) is shown in Figure 12. The important thing to note is that the representation captures the fact that the question is an enquiry about the reference of the truth value of a proposition. The ability to represent this concept supplies the groundwork to build processes (e.g., programs, production rules, inference rules) in the system which uses this concept. The representation can trigger inference rules which will ensure that the system appropriately understands word senses like "ask." "Ask" means something like "enquire into the identity of."

The indirect speech act literature contains many examples of questions such as, "Can you pass the salt?" which are not interpreted in natural discourse as questions. Rather they effect the listener as if he had heard some other illocutionary act; in this case it is a request. However, humans are able to understand that, literally, this sentence is a yes-no question and it follows that they must be able to represent this literal interpretation. Furthermore, most accounts of indirect speech acts build on the assumption that the propositional content of the speech act has been extracted. Hence, we would like to represent this literalness in our network.

We will discuss wh-questions in the section after next section, after we discuss the process of node splitting.

MERGING AND SPLITTING NODES

Merging Nodes

Many writers point out the need for merging nodes which refer to the same physical object (e.g., Anderson, 1977, 1978; Hofstadter, Clossman, & Meredith, 1980). This kind of process seems attractive because it would serve to reduce "clutter" and simplify the crossreferencing problems between information in memory. Common sense seems to dictate that a system which maintains multiple nodes that are known to be co-referential is inefficient. This view is expressed by Hofstadter et al.:

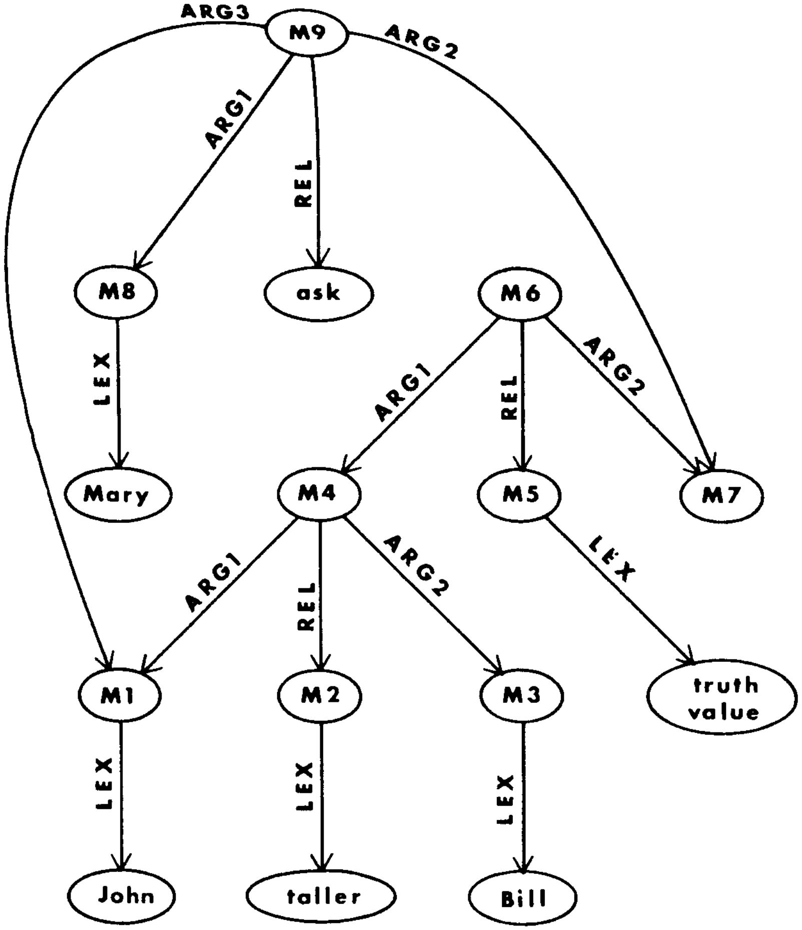

Figure 12. The representation (node M9) for the propositional information contained in the sentence "Mary asks John whether he is taller than Bill."

> It is inevitable in any representational formalism that distinct nodes will sometimes be manufactured that, it turns out later, stand for the same thing. . . . What happens when one finally finds out that the two nodes represent the same thing? Clearly, they have to be fused somehow into one new node. (p. 4)

It is true that this kind of heuristic leads to increased ability to perform inferences in transparent contexts, however it has the disadvantage of not being able to inhibit spurious inference processes in domains involving opaque contexts. For instance, as already pointed out, if a system has two separate nodes for Mary's phone number and Mike's phone number, along with an assertion that they are extensionally equivalent, then it is not difficult to inhibit the system from concluding that John knows Mike's phone number if it is told that John knows Mary's phone number. But if a merge or copy process were to operate on the two nodes known to be extensionally equivalent, as Hofstadter et al., and Anderson propose, so that all of the information separately associated with each of the nodes is merged onto one node, then the spurious inference in the above example would be difficult to inhibit if the system were still going to make transparent inferences. Thus the process of merging nodes must not always take place.

Constraints on Merging Nodes

One possibility for policing this process is to set the constraint that a maximum of one predicate (assertions about co-reference are exempt) can be asserted per node. This amounts to a total suppression of any process which merges nodes; and, although it guarantees a representation with sufficient resolution to appropriately trigger both opaque and transparent inference processes, it creates a maximally inefficient system for cross referencing properties among extensionally equivalent concepts in transparent contexts. This option seems drastic and cumbersome, and we concur with Anderson and Hofstadter et al., in that it does not agree with our intuition either. Furthermore, there is experimental evidence that some merging of nodes does take place in humans. Anderson (1977), in addition to demonstrating that a subject can have two distinct nodes which are extensionally equivalent, has obtained evidence that subjects are strongly biased against maintaining two such distinct nodes if they know that the nodes are extensionally equivalent. Usually in such situations one of the two nodes is less strongly memorized than the other node. What subjects are inclined to do during the course of use is to place duplicate copies of the information which reside at the less established node to the more established node and then gradually abandon use of the less established node.

However, back in McCarthy's telephone number problem, the nodes for Mary's phone number and Mike's phone number cannot be merged. The nodes exist in an opaque context. In all of Anderson's experiments, his stimulus materials were instances of transparent reference. As we said earlier, in transparent contexts, merging of nodes leads to optimal performance.

Let us look at Anderson's (1978) stimulus materials. Subjects were asked to memorize the following kinds of sentences:

(17) The smart Russian is the tall lawyer.
(18) The smart Russian cursed the salesgirl.
(19) The smart Russian rescued the kitten.
(20) The tall lawyer adopted the child.
(21) The tall lawyer caused the accident.

These materials contain no opaque operators. However, we can prefix each of these study sentences with the phrase "John believes that"; then if a subject memorizes these modified study sentences, he or she will generate the same subconfiguration of nodes in memory as before but this time within an opaque context. If a subject fully processes the semantics of the word "believe" (unfortunately, tuning the task demands to force the subject to fully process "believe" might not be so easy) then during the experiment, he or she must process the nodes in a way that maintains the separate intensions. Specifically, the trend in the reaction time data which was observed that indicated subjects were copying information from one co-referent node to another should be absent.

We suggest that the process of merging nodes is inhibited when an opaque context is created but not otherwise. This creates a compromise system which will process instances of transparent reference efficiently yet process instances of opaque reference accurately.

Confusions

We as yet have not succeeded in devising a representation which does not make any semantic confusions. The problem can be seen by examining Figure 2. After the system is told the information in sentence sequence (8), it will answer "yes" to the question "Does Pat know Ed dials Mike's phone number?" The reason is that any straight-forward algorithm that adds the assertion to the data base that Pat knows M4 will also cause it to add to the data base that Pat knows M15 because both nodes have the same status in the network. Humans are certainly able to avoid making this mistake, but the representation can be patched up, either by augmenting the processes which use the representation or by augmenting the representation itself.

Without changing the representation, the only way to inhibit adding the fact that Pat knows node M15 while still allowing the inference that Pat knows node M4 is by using highly domain specific inference rules. Perhaps humans can generate these hypothesized rules quite quickly. It seems that this approach however would lead to a system whose performance was likely to degrade very time it learned something new because there would be the possibility of a new confusion. With each new confusion, a new domain specific inference rule would have to be created to inhibit the confusion. This is reminiscent of an EPAM-like theory of forgetting.

The networks used in this paper use only assertional information. Instead of adding domain specific inference rules, we could augment the representation to include structural, or definitional, information, and then change the semantics of the network so that the meaning of a node was only what the structural information said what it was. Notationally this could be done by adding functional individuals to the present representation constructs (cf. Maida, 1982). The purpose of definitional information would be to define intensions. Although Woods (1975, p. 58) argued for the need to distinguish between structural and assertional information, psychologists have not as yet been inclined to make such a distinction and tend to accept Quillian's formulation (cf. Collins & Loftus, 1975) of a node's full concept as adequate. Anderson (1976, p. 243) does use structural information to represent transparent versus opaque reference but he does not refer to the distinction in his experimental work (e.g., Anderson, 1978).

Splitting

Consider the following situation. What if a memory node created in a transparent context and with several descriptions fused or merged onto it, gets puts into an opaque context? Returning to the phone number problem, what if a cognitive agent learns that Mike and Mary have the same phone number as a consequence of learning that they live together? Then the two descriptions would be attached to the same node as shown in Figure 13. Now suppose this cognitive agent hears sentence (22).

(22) Pete doesn't know that Mike's phone number is the same as Mary's phone number.

The network structure depicted in Figure 13 could not become part of the representation structure for sentence (22) because that structure treats the concepts of Mike's and Mary's phone numbers as being the same, rather than being extensionally equivalent. Our cognitive agent must make a new distinction before he can comprehend (22). It seems necessary that two nodes be constructed to separately represent Mike's phone number and Mary's phone number, and they must be made extensionally equivalent to the original node. We call this process *splitting* and it was mentioned earlier in our discussion of Scott writing *Waverly*. The final representaiton is shown in Figure 14 in which node M16 represents the proposition of (22). This analysis predicts that a person who learns about Mike's and Mary's phone number in a transparent context will have more difficulty subsequently comprehending (22) than a person who learns this same information in an opaque context because a split must take place.

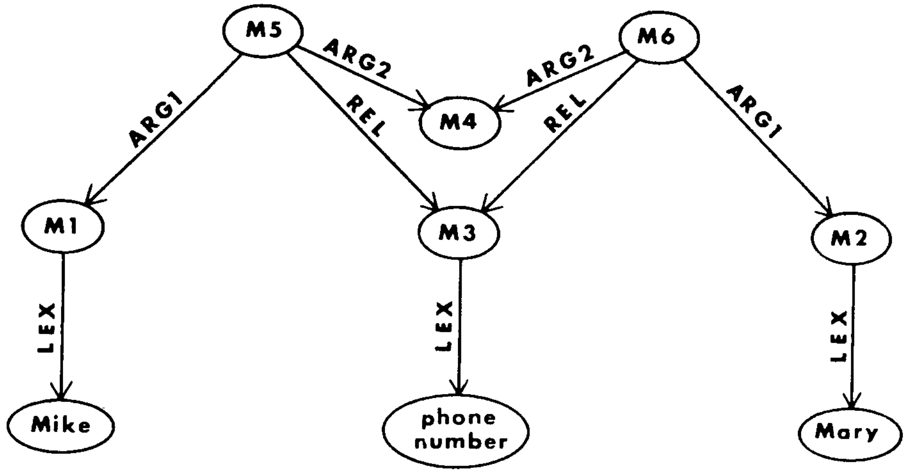

Figure 13. The representation for Mike's and Mary's phone number being the same if there is no immediate need to treat them as separate intensions.

Because of the probable computational overhead involved in splitting nodes, perhaps this kind of strategy might only be utilized when the system is operating in "careful node," as in social situations where misunderstandings could be embarrassing or costly. Also from a cognitive developmental perspective, this operation should occur late in the developmental sequence. And perhaps the usual mode of human functioning leads to treating opaque operators as transparent. Witness the difficulty a teacher has in trying to understand why a student does not understand. If the teacher was facile at processing opaque contexts, he or she would be able to accurately assess the student's knowledge state.

In light of the optional nature of processing associated with opacity, some unanswered experimental questions are raised? What are the factors which trigger opaque processing in humans? Developmentally, when does the ability for opaque processing appear in children? For instance, suppose that a person knows that Jim's wife is Sally's mother and is then told that Mike wants to meet Jim's wife (adapted from Creary, 1979). Under what circumstances might the person *assume that* Mike realizes he would like to meet Sally's mother and under what circumstances might that person feel the need to *determine whether* Mike knows Jim's wife is Sally's mother.

Consider the following two sentences taken from Anderson (1978). Anderson points out that the former sentence is an instance of

(23) I am looking for the best lawyer in town.
(24) I am looking for my little old mother.

opaque reference and the latter is an instance of transparent reference. Given that the difference in interpretation is probably not attributable to their very slight differing syntactic structure, what is determining whether the interpretation is transparent or opaque? The type of interpretation seems to depend on inference processes taking place in the listener which actively assess the knowledge state of the speaker. To illustrate, consider the interpretation of sentence (25) below, depending on whether a foster child, or an earthquake survivor is speaking.

(25) I am looking for my lost mother.

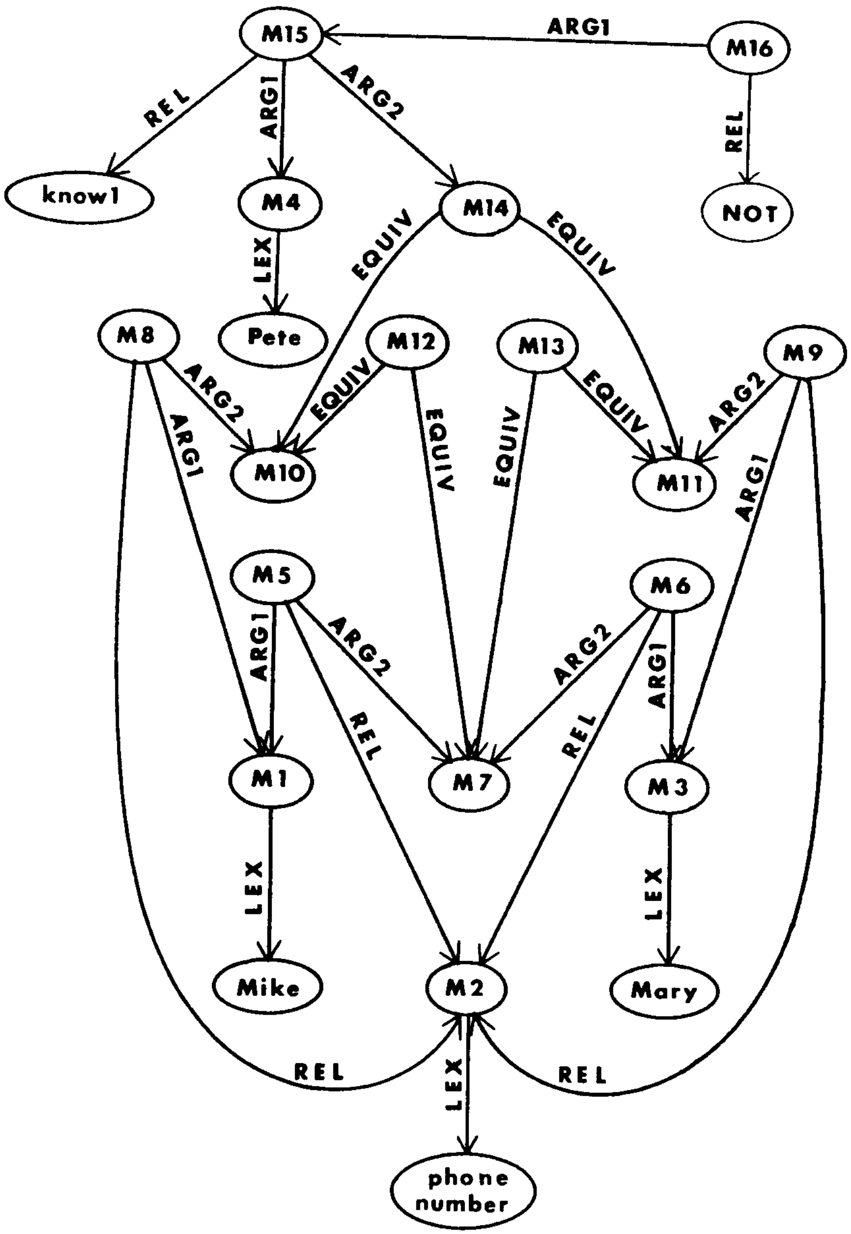

Figure 14. An illustration of the result of a splitting operation.

In the case of the foster child, the listener realizes that the child has probably never known its mother. Whereas in the case of the earthquake survivor, the listener knows that the survivor has someone specific in mind.

REPRESENTING WH-QUESTIONS

The treatment of wh-questions becomes clearer when one recognizes the opacity of the complement position of *ask*. The meaning of sentence (26) is changed if we substitute the phase "Mary's phone number" in place of "Mike's phone number" even if Mike and Mary have the same phone number.

(26) Pat asks Joe for Mike's phone number.

Therefore, in order to represent (26), we must create a new node which represents the intension of the Mike's phone number which Pat asks Joe for. We then equivalence it to the existing node. The resulting interpretation is depicted in Figure 15.

There are additional constraints to consider. We do not want to create a split to answer a simple question if we have already postulated splitting to be computationally expensive; and we want to use as much shared network structure as possible. The representation is constructed and integrated into the rest of memory as follows. The relation *ask* has three arguments: 1) the person doing the asking; 2) the person asked; and, 3) the thing being asked about. Only the thing being asked about is treated opaquely, so it is the only node that does not get matched to the network as the question is being parsed. Rather, it must be asserted to be extensionally equivalent to some other node in memory. Referring to Figure 15, node M8 gets matched to memory via the network pattern determined by nodes M9, M2, and M1. The node which is found as a result of this pattern match (M7 in this case) will be asserted to be extensionally equivalent to node M8. The system should be able to use information which is asserted about M7 to answer the question. In summary then, the listener constructs a representation of the question, matches it to the rest of the network, and finds a node which plays two roles: 1) it is treated as being extensionally equivalent to the concept being asked about; and, 2) it should serve as the access route to information suitable for answering the question. Note also that the system remembers who asked whom what question.

The question is, in fact, an enquiry about a particular intension, and that is the feature the representation unambiguously captures. But it is not so clear for the system as to how to determine what constitutes a good answer to the enquiry about the intension. *This sort of task requires domain specific knowledge. It seems necessary to include discourse rules in the*

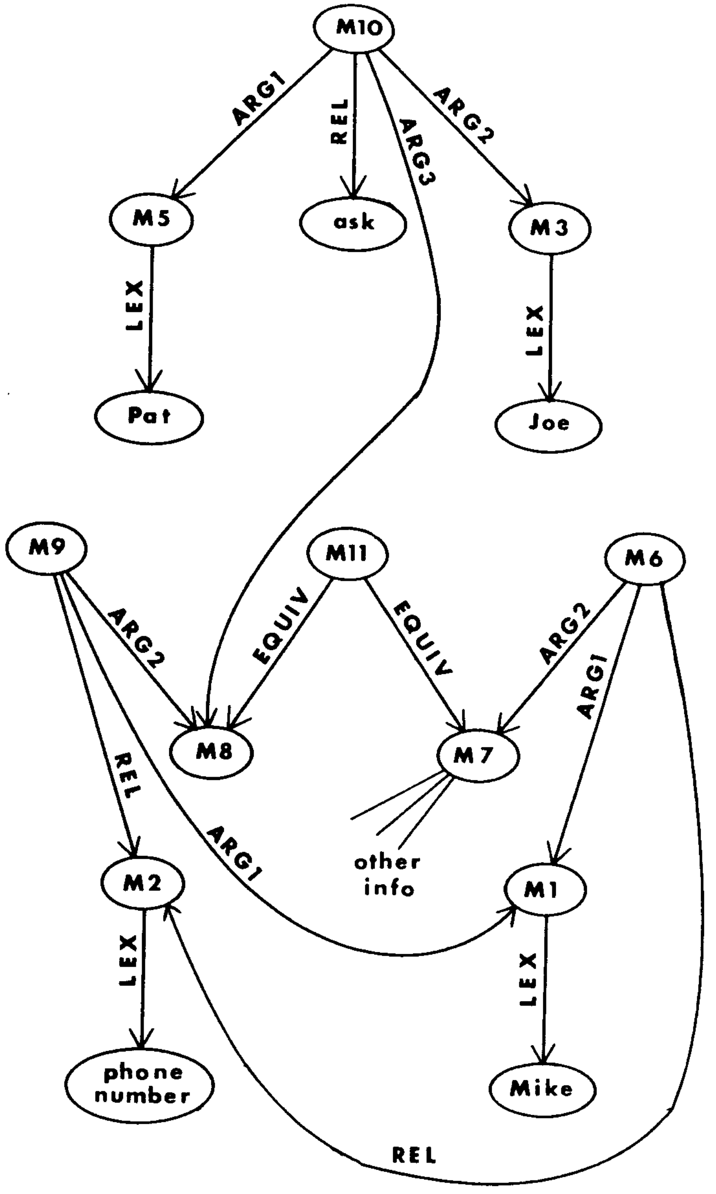

Figure 15. The representation (node M10) for the propositional information contained in the sentence "Pat asks Joe for Mike's phone number."

system to decide exactly what information the questioner is looking for. It is a mysterious phenomenon that humans perceive wh-questions as asking for specific information; on a literal level, they are vague, (e.g., Where is the salt? In the salt shaker). One possible rule for answering questions about phone numbers is given below in (27).

> (27) If a person asks you for a description of a phone number that you believe2 or know2, then he wants you to generate a linguistic description of a digit sequence that is extensionally equivalent to the phone number.

Lehnert (1978), Allen (1979), and Allen and Perrault (1980) have extensively treated the problem of deciding what is so specific about specific questions in natural language discourse.

A more general rule for answering wh-questions, probably reserved for use when the system cannot find relevant domain specific knowledge, is given below in (28).

> (28) When answering a wh-question, use any identifying description of the entity that the questioner asked about, but which the questioner himself did not use in formulating the question.

Thus if Joe answered (26) by saying, "Oh, Mike's phone number" then rule (28) would be violated. Alternatively, if he said "Oh, that's Mary's phone number," then rule (28) would be satisfied but not rule (27).

CONCLUSION

This paper makes an integrated statement concerning the semantic status of nodes in a propositional semantic network, claiming that nodes in such a network represent only intensions. Within the network, the only reference to extensionality should be via some mechanism to assert that two intensions have the same extension (e.g., The Morning Star is the Evening Star). We have also shown how processes which simulate referential transparency or referential opacity can operate on the network. Our analyses have been influenced by analytic philosophy, artificial intelligence, and cognitive psychology.

We have employed this framework in three application areas to illustrate the nature of its solutions. First, we map out a solution to McCarthy's telephone number problem, which he devised to succinctly capture the difference between referential transparency and opacity. Second, we directly represent the notion of the truth value of a proposition as is needed to represent a sentence like "John knows whether he is taller than Bill." Finally, we represent sentences like "Since John asked me whether I was taller than Bill, I assumed he never met Bill."

We also discuss some of the psychological implications of network processes which we call node *merging* and node *splitting*. It has been well recognized that a merge process should occur for co-referential nodes but little attention has been given to restrictions on this merge process and to the need for a splitting process. We theorize that merging should be inhibited in opaque contexts and suggest that Anderson's results on node merging will not generalize to opaque contexts. We also point out the need for a splitting process.

The formalism we use in this paper employs only assertional information and no structural, or definitional, information. This is consistent with network models in the psychological literature. In particular, neither Anderson (1978), nor Collins and Loftus (1975) mention the distinction. Our networks our prone to memory confusions to which the psychologically based networks will also be vulnerable. The notation employed in this paper can be augmented, by the use of functional individuals (cf. Maida, 1982) to eliminate the tendency for memory confusions. Whether psychologists should employ the use of functional individuals is an empirical question. This change in formalism, however, results in a drastic change in the semantic interpretation of network structure, as Quillian proposed it.

Propositional semantic networks continue to be a fruitful data structure by which to model the organization of the human mind.

REFERENCES

Allen, J. F. *A plan-based approach to speech act recognition* (Tech. Rep. 131). Toronto: University of Toronto, Department of Computer Science, 1979.

Allen, J. F., & Perrault, C. R. Analyzing intention in utterances. *Artificial Intelligence*, 1980, *15*, 143-178.

Anderson, J. R. *Language, memory, and thought*. Hillsdale, NJ: Lawrence Erlbaum Associates, 1976.

Anderson, J. R. Memory for information about individuals. *Memory and Cognition*, 1977, *5*, 430-442.

Anderson, J. R. The processing of referring expressions within a semantic network. *Theoretical issues in natural language processing-2*. New York: Association for Computing Machinery, 1978.

Anderson, J. R., & Hastie, R. Individuation and reference in memory: Proper names and definite descriptions. *Cognitive Psychology*, 1974, *6*, 495-541.

Brachman, R. J. What's in a concept: structural foundations for semantic networks. *International Journal for Man-Machine Studies*, 1977, *9*, 127-152.

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

Brachman, R. J., & Smith, B. R. Special issue on knowledge representation. *SIGART Newsletter*, February 1980, (70).

Carnap, R. *Meaning and necessity*. Chicago, IL: University of Chicago Press, 1947.

Cohen, P. R. *On knowing what to say: planning speech acts* (Tech. Rep. 118). Toronto: University of Toronto, Department of Computer Science, 1978.

Cohen, P. R., & Perrault, C. R. Elements of a plan-based theory of speech acts. *Cognitive Science*, 1979, *3*, 177-212.

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

Creary, L. G. Propositional attitudes: Fregean representation and simulative reasoning. *IJCAI Proceedings*, 1979, *6*, 176-181.

Dowty, D. R., Wall, R. E., & Peters, S. *Introduction to Montague Semantics*. Boston, MA: D. Reidel Publishing Company, 1981.

Fahlman, S. *NETL: A system for representing and using real-world knowledge*. Cambridge, MA: MIT Press, 1979.

Frege, G. Über Sinn und Bedeutung. *Zeitschrift für Philosophie und philosophiche Kritik*, 1892, *100*, (all). (Translation available in P. Geach & M. Black (Eds.), *Translations from the philosophical writings of Gottlob Frege*. Oxford: Basil Blackwell, 1960.

Heath, P. L. Nothing. In P. Edwards (Ed.), *The Encyclopedia of Philosophy* (Vol. 5). New York: Macmillan, 1967.

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

Hintikka, J. Semantics for propositional attitudes. In L. Linsky (Ed.), *Reference and modality*. London: Oxford University Press, 1971.

Hofstadter, D. R., Clossman, G. A., & Meredith, M. J. *Shakespeare's plays weren't written by him, but by someone else of the same name* (Tech. Rep. 96). Bloomington, IN: Indiana University, Computer Science Department, 1980.

Kripke, S. Semantical considerations on modal logic. *Acta Philosophica Fennica*, 1963, *16*, 83-89.

Maida, A. S. Using lambda abtraction to encode structural information in semantic networks, 1982, available from the Center for Cognitive Science, Box 1911, Brown University, Providence, RI 02912.

Martin, W. A. Roles, co-descriptors, and the formal representation of quantified English expressions. *American Journal of Computational Linguistics*, 1981, *7*, 137-147.

McCarthy, J. Epistemological problems of artificial intelligence. *IJCAI Proceedings*, 1977, *5*, 1038-1044.

McCarthy, J. First-order theories of individual concepts and propositions. In J. Hayes, D. Michie, & L. Mikulich (Eds.), *Machine Intelligence 9*. New York: Halsted Press, 1979.

McCarthy, J., & Hayes, P. J. Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer & D. Michie (Eds.), *Machine Intelligence 4*. Edinburgh: Edinburgh University Press, 1969.

McCawley, J. Paper given at the Symposition in Cognitive Science, Vassar College, Poughkeepsie, NY, April 25-26, 1980.

McNabb, S. D. *The effects of encoding strategies and age on the memory representation for sentences containing proper names and definite descriptions* (Rep. No. 77-3). Bloomington, IN: Indiana University, Indiana Mathematical Psychology Program, 1977.

Moore, R. C. Reasoning about knowledge and action. *IJCAI Proceedings*, 1977, *5*, 223-227.

Moore, R. C. *Reasoning about knowledge and action* (Tech. Note 191). SRI International, 1980.

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

Quine, W. V. O. Quantifiers and propositional attitudes. *Journal of Philosophy*, 1956, *53*, 177-187.

Russell, B. On denoting. *Mind,* 1906, *14.* (Reprinted in H. Feigl & W. Sellars (Eds.), *Readings in philosophical analysis.* New York: Appleton-Century-Crofts, 1949.

Schank, R. C. *Conceptual information processing.* New York: American Elsevier, 1975.

Schubert, L. K., Goebel, R. G., & Cercone, N. J. The structure and organization of a semantic net for comprehension and inference. In N. Findler (Ed.), *Associative Networks: Representation and use of knowledge by computers.* New York: Academic Press, 1979.

Shapiro, S. C. A net structure for semantic information storage, deduction and retrieval. *IJCAI Proceedings,* 1971, *2,* 512-523.

Shapiro, S. C. Generation as parsing from a network into a linear string. *American Journal of Computational Linguistics,* 1975, Microfiche 33, 45-62.

Shapiro, S. C. Path-based and node-based inference in semantic networks. In D. Waltz (Ed.), *Theoretical issues in natural language processing-2.* New York: Association for Computing Machinery, 1978.

Shapiro, S. C. The SNePS semantic network processing system. In N. Findler (Ed.), *Associative Networks: Representation and use of knowledge by computers.* New York: Academic Press, 1979. (a)

Shapiro, S. C. Generalized augmented transition network grammars for generation from semantic networks. *Proceedings of the 17th Annual Meeting of the Association for Computational Linguistics,* 1979. (b)

Shapiro, S. C., & Neal, J. A knowledge engineering approach to natural language understanding. *Proceedings of the 20th Annual Meeting of the Association for Computational Linguistics,* 1982.

Woods, W. A. What's in a link: the semantics of semantic networks. In D. G. Bobrow & A. M. Collins (Eds.), *Representation and Understanding.* New York: Academic Press, 1975.

APPENDIX

This appendix contains a list of inference rules whose purpose is to propogate assertions across extensional equivalence paths. This is not a complete list to enable all valid inferences, and the specific format of these rules would depend on the specific notation one used. The rules consist of condition-action pairs in which the condition specifies a pattern that must match the network in order for the rule to apply and the action builds more network structure. The condition of a rule consists of a sequence of clauses which all must match the network. Clauses specify network structure as follows:

(1st-arg-slot relation-slot 2nd-arg slot)
or,
(property-slot arg-slot)
or,
(arg-slot equivalent arg-slot)

Variables are prefixed with a question mark and while unbound will match any node in the network provided the rest of the clause matches that part of the network. Once a variable matches a node, it is bound to that node for the rest of the application of that rule.

Rule 1:

```
(is-transparent ?rel) &
(?agent ?rel ?obj1) &
(?obj1 equivalent ?obj2) ==>   (?agent ?rel ?obj2)
```

Note: For example,

If likes is a transparent relation and
John likes Mary and
Mary is Sally's mother then
John likes Sally's mother.

However, he may not believe he likes Sally's mother.

Rule 2:

```
(?agent believe (?agent ?rel ?obj1)) &
(?agent believe (?obj1 equivalent ?obj2))
        ==>   (?agent believe (?agent ?rel ?obj2))
```

If a person believes Jim's wife is Sally's mother and he believes he wants to meet Jim's wife, then he believes he wants to meet Sally's mother regardless of whether Jim's wife in fact is Sally's mother.

Rule 3:

```
(?agent know (?agent ?rel ?obj1)) &
(?agent believe (?obj1 equivalent ?obj2))
        ==>   (?agent believe (?agent ?rel ?obj2))
```

Note: For example,

If John knows he likes Mary and
John believes that Mary is Sally's mother
then he believes he likes Sally's mother.

Although the premises of this rule are stronger than those of rule 2, the conclusion is the same as rule 2.

Rule 4:

```
(?agent know (?agent ?rel ?obj1)) &
(?agent know (?obj1 equivalent ?obj2))
        ==>   (?agent know (?agent ?rel ?obj1))
```

The conclusion of this rule is stronger than that of rule 3 because the system believes the agent is correct in his belief about the coreferentiality of objects 1 and 2.