What do Semantic Network Nodes Represent?

Stuart C. Shapiro

SNERG

Technical Note No. 7

Department of Computer Science

State University of New York at Buffalo

Amherst, New York 14226

Presented at

Foundational Threads in Natural Language Processing

SUNY at Stony Brook

July 27-29, 1981

The work presented here was supported in part by the National Science Foundation under Grant No. MCS80-06314.

Introduction

A semantic network is a representation of knowledge consisting of nodes and labelled, directed arcs in which the following conditions hold: 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 relation 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 question is: what, more precisely, is the nature of a concept represented by a node of a semantic network?

In a recent paper [Maida & Shapiro, 1981], Maida and I argued that nodes of a semantic network represent only intensional concepts. The need for intensional representation follows from the Uniqueness Principle and from the need to represent different knowledge about different intensional concepts that, nevertheless, have the same extensional referent (e.g. the morning star and the evening star). That nodes only represent intensional concepts follows from the difficulty a knowledge representation system would have keeping up with the changing real-world facts about extensional objects. Maida and I showed that taking intensional representation seriously leads to nice solutions to several problems in knowledge representation: McCarthy's "telephone problem"; representing propositions such as "John knows <u>whether</u> he is taller than Bill"; and representing beliefs about questions, e.g. "John asked me whether he was taller than Bill".

In the rest of this paper, I will review the arguments of [Maida & Shapiro, 1981], and try to clarify how an intensional representation can have any connection with the outside world. For this discussion, we will have to consider the input/output behavior of a system using such a representation. We can contemplate two such systems. One is limited to using language. Its input is limited to accepting typed utterances. Its output is limited to typing out expressions of its beliefs. The other system has additional sense organs, such as sight and feeling, and additional effectors, such as hands and feet. These systems must have connections between their conceptual nodes and their sense and effector organs. For the language system, these connections amount to a parser and a generator. These connections can also be represented as knowledge (better, beliefs) of the system. As such, these beliefs form part of the system's understanding of the concepts so connected. Each concept is still intensional, although part of the system's understanding of it may be what it looks like, what it feels like, and how to refer to it in English.

What does a Semantic Network Model?

The first question we must be clear about is what a semantic network as a whole models. One possibility is the real world. This would make nodes represent objects in the world and facts (true assertions) about such objects. Although some people might be interested in semantic networks as models of the world, I am not. I admit that if a semantic network were a model of the world, nodes would have to represent extensional objects. What remains to be discussed is what else a semantic network could model, how that determines an intensional representation, and what connection such a model could have with the world.

The 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, etc. This seems to be the view of some semantic network researchers. In [Woods, 1975], it is argued that the semantic network representation of "The dog that was rabid 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 was rabid." Again, although I feel that semantic networks can, perhaps, be used this way, and this discussion will be resumed in a later section, this is not the main use to which I would put semantic networks.

A third possibility, and the one I prefer, is that a semantic network models the knowledge 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.

The Need for Intensional Representation

The need for the notion of an intension or sense, as opposed to an extension or reference, was first pointed out in [Frege, 1892] to explain the difference between the sentences "The morning star is the evening star" and "The morning star is the morning star." The first sentence states a fact that might be news to somebody, the second is a tautology and provides no information. The phrases "the morning star," and "the evening star" have the same extension, but different intensions. The significance of this to the semantics of semantic networks was discussed in [Woods, 1975]. If nodes of a semantic network represented only extensions, the Uniqueness Principle would require that the node representing the morning star be the same as the node representing the evening star, and the proposition that the morning star is the evening star could not be distinguished from the proposition that the morning star is the morning star.

Woods concluded that some, but not all, nodes of a semantic network had to represent intensions. Brachman [1977] put forth the argument that "Semantic network nodes are representations of the intensions of natural language designators" [Brachman, 1977, p. 139, italics in the original], yet "some of the operations in the network scheme are purely extensional" [Brachman, 1977, p. 150]. In [Maida & Shapiro, 1981], it is argued that all nodes of a semantic network represent only intensions.

Some Implications of Intensional Representation

The Need for Co-referential Propositions

If a semantic network has a node for <the morning star> (I will denote an intension by enclosing a designating phrase in angle brackets) and a different node for <the evening star>, what should be done when the assertion is made that the morning star is the evening star? The wrong solution is to merge the two nodes by transferring all arcs from one node to the other and eliminating the first. This would eliminate the distinction of the two intensional concepts and make it impossible to represent "John did not know that the morning star is the evening star" differently from "John did not know that the morning star is the morning star." The proper solution is to add a node to the network representing the proposition, <the morning star is co-referential with the evening star>. This maintains two nodes for the two distinct intensions, while adding the proposition that they are co-referential. The co-referential proposition can be used by the system's reasoning processes to infer that certain beliefs about <the morning star> can be transferred to <the evening star>. This will be discussed further below.

Transient-Process Account

In [Woods, 1975], the "transient-process account" is criticized. This account says that restrictive relative clauses do not, normally, add information to the network, but are used to find an old node which then gets new information added to it from the main clause. However, if we view a semantic network as a model of a conceptual memory, and if we follow the Uniqueness Principle, we must accept the transient process account. If the system is told that a dog has rabies, and then is told, "The dog that has rabies bit a man," the system can only be said to have understood that sentence if it adds to its intensional concept of the rabid dog the information that it bit a man (assuming that it believes what it is told). If it builds a new node for the rabid dog of the latter sentence, the Uniqueness Principle (UP) tells us that it is a different dog from the first (at least intensionally). Even if the same dog-node is used, but a new node is created for the subordinate clause proposition that it has rabies, this would be in conflict with the UP, because there would be two nodes for the same proposition (an intensional concept), that this particular dog has rabies. If the transient process account is followed, the system is left with one dog, about which it believes that it has rabies and that it bit a man. This seems to be the right situation.

Order Dependency

The above two sections imply that if a semantic network based understanding system learns information about two concepts, then learns that they are co-referential, the network will have a different set of nodes than if it first learned that the two concepts were coreferential. For example, after learning "Scott wrote Ivanhoe" and "The author of Ivanhoe wrote The Lady of the Lake", there will be only one node for <Scott, the author of Ivanhoe and The Lady of the Lake>. However, if the system now learns that someone wrote Waverly, it will have a separate node for <the author of Waverly>, and if it then learns that Scott is the author of Waverly, it will just add a co-referential proposition that <the author of Waverly> is <Scott, the author of Ivanhoe and The Lady of the Lake>. If it had heard with the first mention of Waverly that Scott wrote it, the separate node for <the author of Waverly> would never have been created.

Referential Opacity as the Norm

Referentially opaque contexts seem to be a problem for reasoning systems because extensional representation is assumed, and, therefore, the substitutivity of equals for equals is assumed. An opaque context must somehow block this normal way of reasoning. This is why McCarthy's telephone problem [McCarthy, 1979] is difficult. In this problem, it is known that Mary's telephone number is Mike's telephone number. From "Pat dialled Mike's telephone number" it should follow that "Pat dialled Mary's telephone number," but from "Pat knows Mike's telephone number," it should not follow that "Pat knows Mary's telephone number." In extensional systems, the first inference is easy but it is hard to block the second.

A system that uses the Uniqueness Principle does not need the substitutivity of equals for equals as a basic reasoning rule, because no two distinct nodes represent truely equal concepts. If two propositions are about the same intensional concept, they already share the node for this concept. Co-referentiality is not the same as equality. Co-referentiality is merely another proposition that can be used in reasoning. We can catagorize certain predicates, such as dialling but not knowing, as referentially transparent. This can be done by nodes in the network representing propositions about these predicates. We can also have a reasoning rule, again represented in the network, that says that if a referentially transparent predicate applies to a concept, then it also applies to any co-referential concept. That rule will explicitly allow the dialling inference to go through, but there will be no such rule to permit the knowing inference. Intensional representation implies that referential opacity is the norm and referential transparency must be explicitly sanctioned. Since it is easier to allow something special than to block something which is normal, the telephone problem and similar problems are easily solved.

Existence and Non-existence

Since a node in a semantic network represents an intensional concept, rather than an extensional object, it carries no commitment that a real world object corresponding to the intensional concept exists. A naive 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 -- (THERE-IS-AN x) [x is The-present-king-of-France & x is Bald]. However, this is because of the normal extensional interpretation of statements of standard logic. A constant node in a semantic network is like a Skolem constant derived from an existentially quantified variable that asserts only the existence of the intensional concept. Existence of an object in the world is only another property implying our expectation of being able to see, meet, touch, interact with, etc. the object. "The present king of France is bald" can be represented in a semantic network by a node representing the proposition that the intensional concept <the present king of France> has the property of being bald. That is, there is a node representing the concept, <the bald, present king of France>. Nothing has yet been stored about whether we can expect to meet up with an object matching this concept in the real world. The statement, "The present king of France does not exist" calls to mind (creates in the conceptual belief structure) the intensional concept of <the present king of France> and asserts about him that neither we nor anyone we know nor anyone we know of can expect to meet such a person. Nevertheless, <the present king of France> is connected in the network to our concepts of kingship, of France, and of the present time, and so a lot is known about the concept and what an object matching it would be like.

Just as a node can represent <the present king of France> without commitment to its objective reality, nodes can represent theoretical concepts that are useful for understanding the world without commitment to their objective reality. For example, there can be nodes representing <short term memory>, <electrons>, <gluons>, even <propositions>, <intensional concepts>, and <truth>.

Connections with Reality

The main objection to exclusive intensional representation seems to be that if nodes only represent 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. Since so many AI understanding systems deal exclusively with language, rather than with sight, manipulators, legs or wheels, etc., we can consider two different systems. One is a language system. Its sense organ is a keybord. Its only effector organ is a CRT screen. The other system is a robot with sight and manipulators as well as language.

Since the language system interacts with the outside world only through language, the only questions we can consider about the connections of its concepts with reality are questions such as "Does it use words like we do?", "When it uses word ________, does it mean the same thing as when I use it?" and "When I use word ________, does it understand what I mean?"

The perceptual system of the language system is its parser/analyzer -- the programs that analyze typed utterances and build peices of semantic network. The motor system is the generator -- the programs that analyze a section of the semantic network and construct an utterance to be displayed on the CRT. One crucial requirement for an adequate connection with the world is simple consistency of input/output behavior. That is, a phrase that is analyzed to refer to a particular node should consistently refer to that node at least while there is no change in the network. Similarly, if the system generates a certain phrase to describe the concept represented by a node, it should be capable of generating that same phrase for that some node while nothing in the network changes. Notice that it is unreasonable to require that if a phrase is generated to describe a node, the analyzer should be able to find the node from the phrase. The system might know of several brown dogs and describe one as "a brown dog". It could not be expected to find that node as the representation of "a brown dog" consistently.

Much of the analyzer and the generator can be put in the network itself. A prototype system is described in [Neal, 1981] that represents words and sentences (i.e. the intensional concepts of words and sentences) in a semantic network, as well as rules that express how language strings relate to the propositions they express. These rules comprise the systems knowledge of how language is used -- what utterances expressed by others to it mean and how it should express its concepts. The perceptual and motor systems "hidden in programs" rather than expressed in the semantic network knowledge structure are very small.

If we are assured of the simple input/output consistency of the system, the main question left is does it use words to mean the same thing as we do. It is the same question we would be concerned with if talking with a blind invalid, although in that case, we would assume the answer was "yes" until the conversation grew so bizzare that we were forced to change our minds. As the system, or the invalid, utters more and more sentences using a particular word or phrase we will become more and more convinced that it means what we would mean by it, it means what we would describe by a different word or phrase ("Oh! When you say "conceptual dependency structure", you mean what I mean when I say "semantic network."), we don't know what is meant or we are convinced that it is not used in a consistent, meaningful way, and the system (or invalid) does not know what it is talking about. As long as the conversation proceeds without our getting into the latter situation, the system has all the connections with reality it needs.

The situation is similar with the robot. It needs a perceptual system in which some node, set of features, etc. is triggered consistently when a given object is seen. If it is to communicate reasonably with people, it must be able to make approximately the same perceptual distinctions that we make, although color blind people can certainly communicate well enough with non-color blind people. These perceptual nodes need not extensionally represent the objects that trigger them. My chair in my office may look exactly like your chair in your office, nevertheless, I can recognize my chair in its usual context and I have knowledge about it. The perceptual nodes can be connected to semantic-conceptual nodes, just like the nodes for words and sentences can be connected to nodes for propositions. 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). Finally, it needs hand-eye coordination, which is comparible to the conversational consistency required by the language system. None of this requires that semantic network nodes be somehow tied to particular extensional objects in the world.

Summary

I have argued that the properties of semantic networks as representations of knowledge, particularly the Uniqueness Principle, imply that semantic network nodes represent intensional concepts exclusively. I reviewed the implications of this view, which include some nice solutions to some problems in the representation of knowledge. Finally, I tried to show that intensional representation does not cause an understanding system so based to be cut off from the world in such a way that its operation is meaningless.

References

1. Brachman, R. J. What's in a concept: structural foundations for semantic networks. Int. J. Man-Machine Studies 9, (1977), 127-152.

2. Frege, G. Uber Sinn und Bedeutung. Zeitschrift fur Philosophie und philosophiche Kratik, 1892, 100 (all).

Translated in P. Geach and M. Black, eds. Translations from the philosophical writings of Gottlob Frege. Oxford: Basil Blackwood, 1960, 56-78.

3. Maida, A. S. and Shapiro, S. C., Intensional Concepts in Propositional Semantic Networks, Technical Report No. 171, Dept. of Computer Science, SUNY at Buffalo, Feb., 1981.

4. McCarthy, J. First order theories of individual concepts and propositions. In J. Hayes, D. Michie, and L. Mikulich, eds. Machine Intelligence 9. New York: Halsted Press, 1979, 129-147.

5. Neal, J. A knowledge engineering approach to natural language understanding, Technical Report No. 179, Department of Computer Science, SUNY at Buffalo, Amherst, NY, June, 1981.

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