proceedings

aaai-86

*fifth national conference on
artificial intelligence*

august 11–15, 1986   philadelphia, pa

*volume 1
science*

SNePS CONSIDERED AS A FULLY INTENSIONAL PROPOSITIONAL SEMANTIC NETWORK

Stuart C. Shapiro and William J. Rapaport

Department of Computer Science
State University of New York at Buffalo
Buffalo, NY 14260
{ rapaport | shapiro }%buffalo@csnet-relay

ABSTRACT

We present a formal syntax and semantics for SNePS considered as the (modeled) mind of a cognitive agent. The semantics is based on a Meinongian theory of the intensional objects of thought that is appropriate for AI considered as "computational philosophy" or "computational psychology".

1. INTRODUCTION.

We present a formal syntax and semantics for the SNePS Semantic Network Processing System (Shapiro 1979), based on a Meinongian theory of the intensional objects of thought (Rapaport 1985a). Such a theory avoids possible worlds and is appropriate for AI considered as "computational philosophy"—AI as the study of how intelligence is possible—or "computational psychology"—AI with the goal of writing programs as models of human cognitive behavior. Recently, SNePS has been used for a variety of AI research and applications projects. These are described in Shapiro & Rapaport 1985, of which the present paper is a much shortened version. Here, we use SNePS to model (or construct) the mind of a cognitive agent, referred to as CASSIE (the Cognitive Agent of the SNePS System—an Intelligent Entity).

2. INTENSIONAL KNOWLEDGE REPRESENTATION.

SNePS represents *propositions* about entities having properties and standing in relations. Nodes represent the propositions, entities, properties, and relations, while the arcs represent structural links between these. SNePS nodes *might* represent *extensional* entities, whose identity conditions do not depend on their manner of representation. Two extensional entities are equivalent (for some purpose) iff they are identical (i.e., iff "they" are really one entity, not two).

Although SNePS *can* be used to represent extensional entities in the world, we believe that it *must* represent *intensional* entities—entities whose identity conditions *do* depend on their manner of representation. Two intensional entities might be equivalent (for some purpose) without being identical (i.e., they might really be two, not one). Only if one wants to represent the relations between a mind and the world would it also have to represent extensional entities (Rapaport 1978, McCarthy 1979). If SNePS is used just to represent a mind—i.e., a mind's model of the world—then *it does not need to represent any extensional objects.* It can then be used either to model the mind of a particular cognitive agent or to build such a mind—i.e., to *be* a cognitive agent itself (Maida & Shapiro 1982). There have been a number of arguments presented in both the AI and philosophical literature in the past few years for the need for intensional entities (Castañeda 1974, Woods 1975, Maida & Shapiro 1982, Rapaport 1985a, Brach-

man 1977, Routley 1979, Parsons 1980). Among them, the following considerations are especially significant:

Principle of Fine-Grained Representation: The objects of thought (i.e., intentional objects) are intensional; a mind can have two or more objects of thought that correspond to only one extensional object. To take the classic example, the Morning Star and the Evening Star might be distinct objects of thought, yet there is only one extensional object (a certain astronomical body) corresponding to them.

Principle of Displacement: Cognitive agents can think and talk about non-existents: a mind can have an object of thought that corresponds to no extensional object. Again to take several classic examples, cognitive agents can think and talk about fictional objects such as Santa Claus, possible but non-existing objects such as a golden mountain, impossible objects such as a round square, and possible but not-yet-proven-to-exist objects such as theoretical entities (e.g., black holes).

If nodes only represent intensions (and extensional entities are not represented in the network), how do they link up to the external, extensional world? One answer is by means of a LEX arc (see (Syn.1) and (Sem.1), below). The nodes at the head of the LEX arc are *our* (the user's) interpretation of the node at its tail. The network without the LEX arcs and their head-nodes displays the *structure* of CASSIE's mind (Carnap 1928, Sect. 14; for other answers, see Maida & Shapiro 1982, Shapiro & Rapaport 1985).

3. DESCRIPTION OF SNePS.

SNePS satisfies the *Uniqueness Principle:* There is a one-to-one correspondence between nodes and represented concepts. This principle guarantees that nodes represent intensional objects and that nodes will be shared whenever possible. Nodes that only have arcs pointing *to* them are considered to be unstructured or *atomic.* They include: (1) *sensory* nodes, which—when SNePS is being used to model a mind—represent interfaces with the external world (in the examples that follow, they represent utterances); (2) *base* nodes, which represent individual concepts and properties; and (3) *variable* nodes, which represent arbitrary individuals (Fine 1983) or arbitrary propositions.

Molecular nodes, which have arcs emanating *from* them, include: (1) *structured individual* nodes, which represent structured individual concepts or properties (i.e., concepts and properties represented in such a way that their internal structure is exhibited); and (2) *structured proposition* nodes, which represent propositions; those with no incoming arcs represent *beliefs* of the system. (Note that structured proposition nodes can also be considered to be structured individuals.) Proposition nodes are either *atomic* (representing atomic propositions) or are *rule nodes.* Rule nodes represent deduction rules and are used for node-based deductive inference (Shapiro 1978; Shapiro & McKay 1980; McKay & Shapiro 1981; Shapiro, Martins, & McKay 1982). For each of the three categories of molecular nodes (structured individuals, atomic propositions, and rules), there are *constant* nodes of that category and *pattern* nodes of that category representing arbitrary entities of that category.

There are a few built-in arc labels, used mostly for rule nodes. *Paths* of arcs can be defined, allowing for *path-based* inference, including property inheritance within generalization hierarchies (see below; Shapiro 1978, Srihari 1981). All other arc labels are defined by the user, typically at the beginning of an interaction with SNePS.

3.1. CASSIE—A Model of a Mind.
Since most arcs are user-defined, users are obligated to provide a formal syntax and semantics for their SNePS networks. We shall describe the way in which we have been using SNePS to build CASSIE. Using Brachman's (1979) terminology, insofar as SNePS is a semantic network system at the logical level and can thus be used to define one at the epistemological or conceptual level, CASSIE is SNePS being used at a conceptual level.

The nodes represent the objects of CASSIE's thoughts—the things she thinks about, the properties and relations with which she characterizes them, her beliefs, etc. (Maida & Shapiro 1982, Rapaport 1985a). The Principle of Displacement says that a cognitive agent is able to think about virtually anything, including fictional objects, possible but non-existing objects, and impossible objects. Any theory that would account for this requires a nonstandard logic, and its semantics cannot be limited to merely *possible* worlds.

Theories based on Alexius Meinong's Theory of Objects are of precisely this kind. Meinong held that psychological experiences consist in part of a psychological *act* (such as thinking, believing, wishing, etc.) and the *object* to which the act is directed (e.g., the object that is thought about or the proposition that is believed). Two kinds of Meinongian objects of thought are relevant for us: (1) The *objectum*, or object of "simple" thoughts: Santa Claus is the objectum of John's act of thinking of Santa Claus. The meaning of a noun phrase is an objectum. (2) The *objective*, or object of belief, knowledge, etc.: that Santa Claus is thin is the objective of John's act of believing that Santa Claus is thin. Objectives are like propositions in that they are the meanings of sentences and other sentential structures. Note that objecta need not exist and that objectives need not be true. (Cf. Meinong 1904; Rapaport 1978, 1981; Castañeda 1974, 1975; Routley 1979; Parsons 1980; Lambert 1983; Zalta 1983.)

This is, perhaps, somewhat arcane terminology for what might seem like AI common sense. But without an underlying theory, such as Meinong's, there is no way to be sure if common sense can be trusted. It is important to note that not only are all represented things intensional, but that they are all objects of CASSIE's mental acts; i.e., they are all in CASSIE's mind (her "belief space")—they are all intentional. Thus, even if CASSIE represents the beliefs of someone else (e.g., John's belief that Lucy is rich, as in the conversation in Sect. 3.2), the objects that she represents as being in that person's mind (as being in his "belief space") are actually CASSIE's representations of those objects—i.e., they are in CASSIE's mind.

3.2. A Conversation with CASSIE.
Before giving the syntax and semantics of the case-frames employed in representing CASSIE's "mind", we present a conversation we had with her, showing the network structure as it is built—i.e., showing the structure of CASSIE's mind as she is given information and as she infers new information. An ATN parser/generator (Shapiro 1982) was used to parse the English input into SNePS and the SNePS structures into English. User input is on lines with the :-prompt; CASSIE's output is on the lines that follow. Comments are enclosed in brackets. A fragment of the full network showing CASSIE's state of mind at the end of the conversation is shown in Fig. 1.

: Young Lucy petted a yellow dog
I understand that young Lucy petted a yellow dog

[CASSIE is told something, which she now believes. At this point, her entire belief structure consists of nodes b1, m1-m13, and the corresponding sensory nodes. The node labeled "now" represents the current time, so the petting is clearly represented as being in the past. CASSIE's response is "I understand that" appended to her English description of the proposition just entered.]

What is yellow
a dog is yellow

[This response shows that CASSIE actually has some beliefs; she did not just parrot back the above sentence.]

: Dogs are animals
I understand that dogs are animals

[CASSIE is told a small section of a class hierarchy.]

: Who petted an animal
young Lucy petted a yellow dog

[CASSIE can answer the question using the class hierarchy, because, prior to the conversation, an inheritance rule was given to SNePS. The rule says that the CLASS arc is implied by the path consisting of a CLASS arc followed by zero or more occurrences of the two-arc path consisting of the converse SUBCLASS arc followed by the SUPERCLASS arc (Shapiro 1978, Srihari 1981). The dog was called "a yellow dog" rather than "a yellow animal" because the redundant CLASS arc is not built. CASSIE's new beliefs about the dog's classification and color consists of nodes m9-m13, m16-m17, and the corresponding sensory nodes.]

: Lucy is sweet
I understand that young Lucy is sweet

[CASSIE's response shows that she identifies this Lucy with the previous Lucy.]

: Lucy is a girl
I understand that sweet young Lucy is a girl

[The beginning of a class hierarchy for Lucy. Notice that all the adjectival properties of Lucy are mentioned.]

: Girls are people
I understand that girls are people

[More of the class hierarchy is given.]

: What was petted by a person
a yellow dog was petted by sweet young Lucy

[The proposition is retrieved using the CLASS inheritance rule.]

: The dog is named Rover
I understand that Rover is a yellow dog

['The dog' refers to the only dog CASSIE knows about, who is now given a name.]

: What was petted by a person
yellow Rover was petted by sweet young Lucy

[This question was asked before but is answered differently this time, because the dog now has a name, and CASSIE prefers to describe an individual by its name when it has one.]

: John believes that Lucy is rich
I understand that John believes that Lucy is rich

[In our current development of CASSIE, she interprets 'believes that' contexts to be *de dicto*, so she assumes that the Lucy that John has in mind is a different one from the Lucy that she knows. CASSIE's beliefs about the two Lucies are added to the network (nodes m40, m44).]
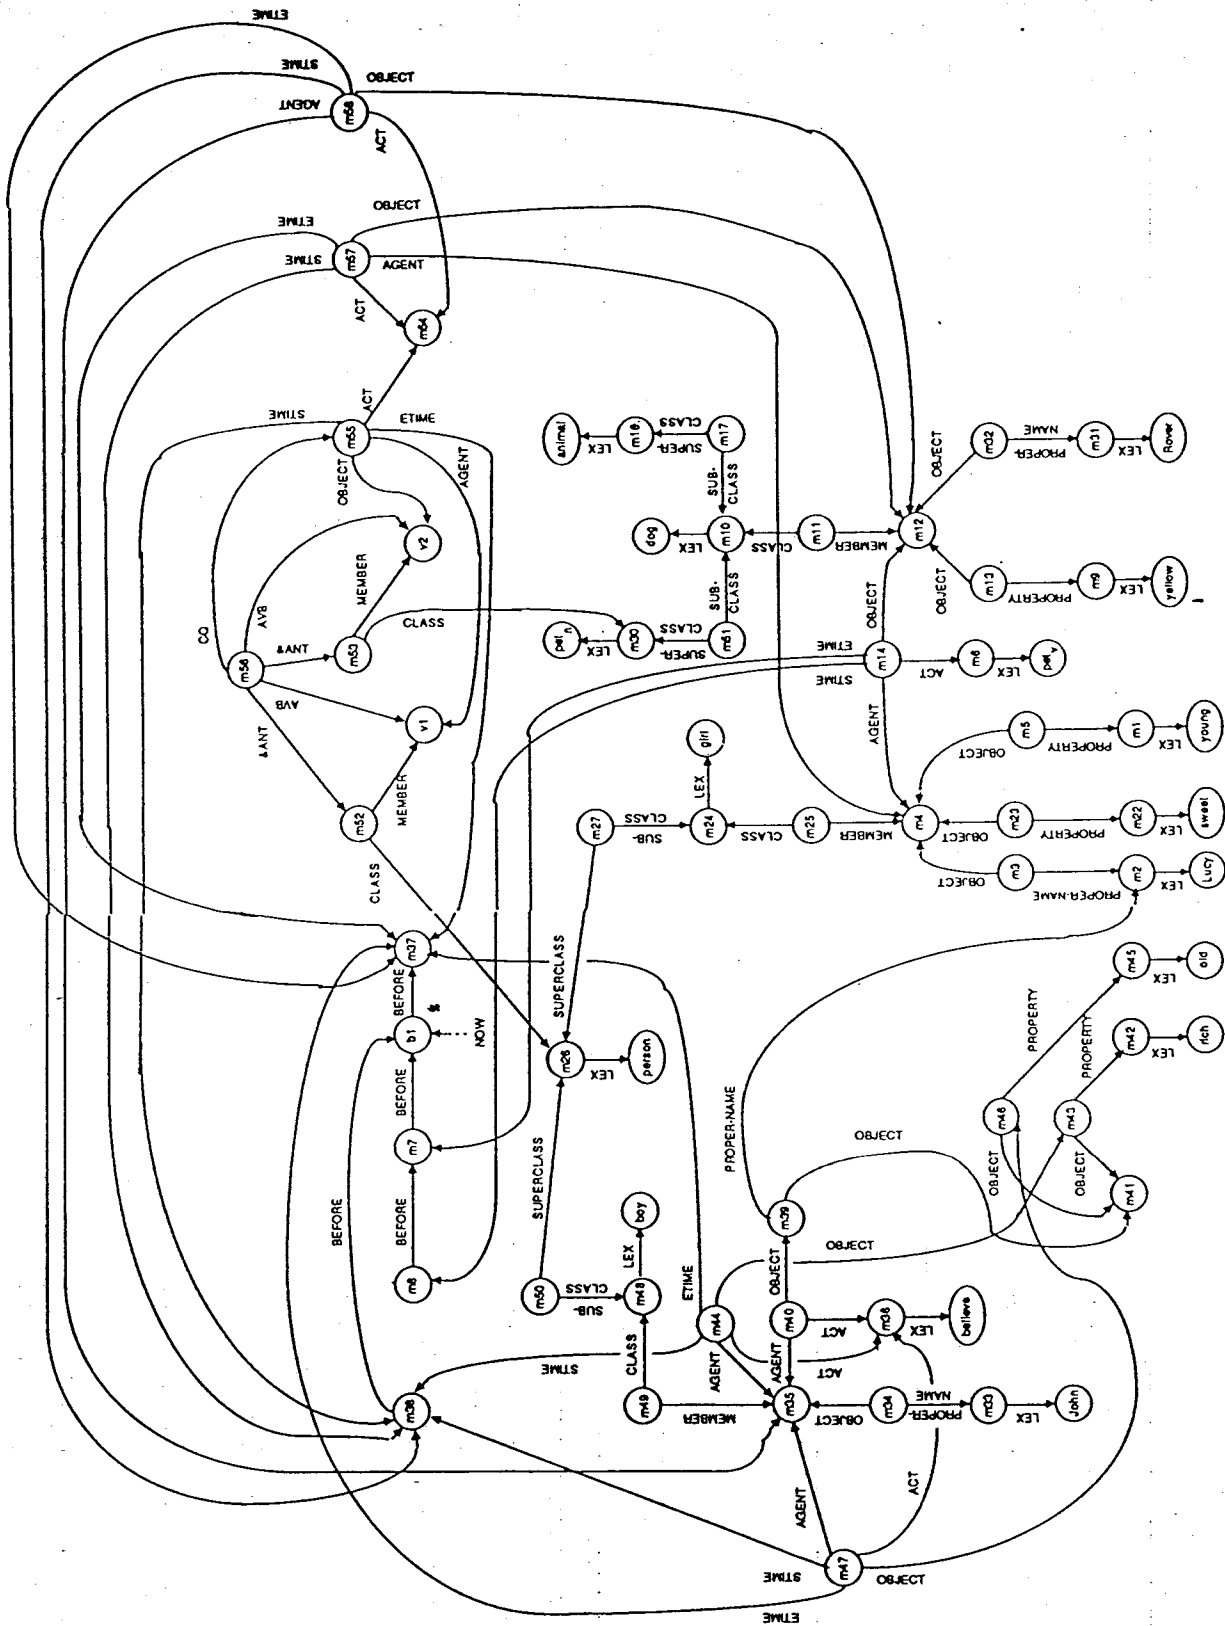
Figure 1.

: Who is rich
I don't know.

[CASSIE knows no one who is rich. She only believes that *John* believes that someone (whom he believes to be named 'Lucy') is rich. The answer is 'I don't know', rather than 'no one is rich', because CASSIE doesn't use the closed-world hypothesis.]

: Who is sweet
young Lucy is sweet

[This question demonstrates that Lucy is able to answer a "who is <property>" question when she has relevant beliefs.]

: John believes that Lucy is old
I understand that John believes that rich Lucy is old

[CASSIE assumes that John knows a different Lucy than she knows, but she assumes that all John's beliefs about "Lucy" are about the same Lucy.]

: John is a boy
I understand that John is a boy

[This and the next two inputs are given to establish more of the class hierarchy and to make it clear that when CASSIE answers the last question of this session, she is doing both path-based reasoning and node-based reasoning at the same time.]

: Boys are people
I understand that boys are people

: Dogs are pets
I understand that dogs are pets

: For every p and d if p is a person and d is a pet then p loves d
I understand that for every d and p, if p is a person and d is a pet then p loves d

[This node-based rule fits into the class hierarchy as node m56. This is, we believe, equivalent to the integrated TBox/ABox mechanism proposed for KRYPTON (Brachman et al. 1983, Brachman et al. 1985).]

: Who loves a pet
sweet young Lucy loves yellow Rover and
John loves yellow Rover
[The question was answered using path-based inferencing to deduce that Lucy and John are people and that Rover is a pet, and node-based inferencing to conclude that, therefore, Lucy and John love Rover.]

3.3. Syntax and Semantics of SNePS.

In this section, we give the syntax and semantics of the nodes and arcs used in the interaction. What we present here is our current model; we make no claims to completeness of the representational scheme. We begin with a few rough definitions. (Cf. Shapiro 1979, Sect. 2.1, for more precise ones.)

(Def. 1) A node *dominates* another node if there is a path of directed arcs from the first node to the second node.

(Def. 2) A *pattern* node is a node that dominates a variable node.

(Def. 3) An *individual* node is either a base node, a variable node, or a structured constant or pattern individual node.

(Def. 4) A *proposition* node is either a structured proposition node or an atomic variable node representing an arbitrary proposition.
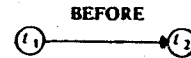
(Syn.1) If "w" is a(n English) word and "i" is an identifier not previously used, then

$$i \xrightarrow{\text{LEX}} w$$
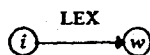
is a network, w is a sensory node, and i is a structured individual node.

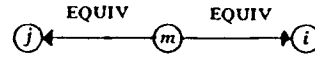(Sem.1) i is the objectum corresponding to the utterance of w.

(Syn.2) If either "t_1" and "t_2" are identifiers not previously used, or "t_1" is an identifier not previously used and t_2 is a temporal node, then

$$t_1 \xrightarrow{\text{BEFORE}} t_2$$

is a network and t_1 and t_2 are *temporal* nodes, i.e. individual nodes representing times.

(Sem.2) t_1 and t_2 are objecta corresponding to two times, the former occurring before the latter.

(Syn.3) If i and j are individual nodes, and "m" is an identifier not previously used, then

$$j \xleftarrow{\text{EQUIV}} m \xrightarrow{\text{EQUIV}} i$$

is a network and m is a structured proposition node.

(Sem.3) m is the objective corresponding to the proposition that objecta i and j (are believed by CASSIE to) correspond to the same actual object. (This is not used in the conversation, but is needed for fully intensional representational systems; cf. Rapaport 1978, 1984b; Castañeda 1974; Maida & Shapiro 1982.)
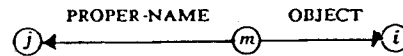
(Syn.4) If i and j are individual nodes and "m" is an identifier not previously used, then

$$j \xleftarrow{\text{PROPERTY}} m \xrightarrow{\text{OBJECT}} i$$
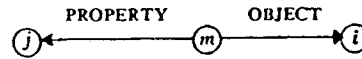
is a network and m is a structured proposition node.

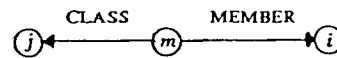(Sem.4) m is the objective corresponding to the proposition that i has the property j.

(Syn.5) If i and j are individual nodes and "m" is an identifier not previously used, then

$$j \xleftarrow{\text{PROPER-NAME}} m \xrightarrow{\text{OBJECT}} i$$

is a network and m is a structured proposition node.

(Sem.5) m is the objective corresponding to the proposition that objectum i's proper name is j. (j is the objectum that is i's proper name; its expression in English is represented by a node at the head of a LEX-arc emanating from j.)
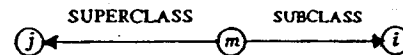
(Syn.6) If i and j are individual nodes and "m" is an identifier not previously used, then

$$j \xleftarrow{\text{CLASS}} m \xrightarrow{\text{MEMBER}} i$$

is a network and m is a structured proposition node.

(Sem.6) m is the objective corresponding to the proposition that i is a (member of class) j.

(Syn.7) If i and j are individual nodes and "m" is an identifier not previously used, then

$$j \xleftarrow{\text{SUPERCLASS}} m \xrightarrow{\text{SUBCLASS}} i$$

is a network and m is a structured proposition node.

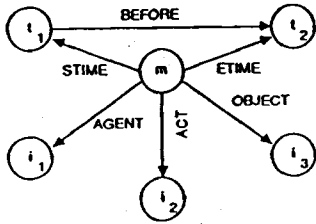(Sem.7) m is the objective corresponding to the proposition that (the class of) is are (a subclass of the class of) js.

(Syn.8) If i_1, i_2, i_3 are individual nodes, t_1, t_2 are temporal nodes, and "m" is an identifier not previously used, then

is a network and *m* is a structured proposition node. (Nodes m40, m44 are examples of this for the mental act of believing; cf. Rapaport 1984b, Rapaport & Shapiro 1984. The ETIME and STIME arcs are optional and can be part of any proposition node; they are a provisional technique for handling temporal information—cf. Shapiro & Rapaport 1985.)

(Sem.8) *m* is the objective corresponding to the proposition that agent i_1 performs act i_2 with respect to i_3 starting at time t_1 and ending at time t_2, where t_1 is before t_2.

Rule nodes have been described more fully in Shapiro 1979, and a full syntax and semantics for them is presented in Shapiro and Rapaport 1985. Here, we present the syntax and semantics only for the node-based inference rule used in the conversation with CASSIE (Fig. 2, node m56):

(Syn.9) If $a_1, \ldots, a_n, c_1, \ldots, c_j$, and $d_1, \ldots, d_k$ are proposition nodes ($n, j, k \geq 0$), and "r" is an identifier not previously used, then
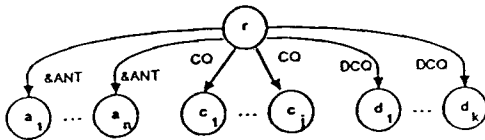
is a network, and *r* is a rule node.

(Sem.9) *r* is the objective corresponding to the proposition that the conjunction of the propositions $a_1, \ldots, a_n$ relevantly implies each c_i ($1 \leq l \leq j$) and relevantly implies each d_i ($1 \leq l \leq k$) for which there is not a better reason to believe it is false. (The d_i are *default* consequences: each is implied only if it is neither the case that CASSIE already believes *not* d_i nor that *not* d_i follows from non-default rules.)

(Syn.10) If *r* is a rule node, and *r* dominates variable nodes $v_1, \ldots, v_n$, and, in addition, arcs labeled "AVB" go from *r* to each v_i, then *r* is a quantified rule node.

(Sem.10) *r* is the objective corresponding to the proposition that the rule that would be expressed by *r* without the AVB arcs holds after replacing each v_i by any object in its range.

4. SNePS AND CASSIE AS SEMANTIC NETWORKS.

We conclude by looking at SNePS from the perspective of Brachman's discussions of structured inheritance networks and hierarchies of semantic-network formalisms (Brachman 1977, 1979). Brachman offers six criteria for semantic networks:

A semantic network must have a *uniform notation*. SNePS provides some uniform notation with its built-in arc labels for rules, and it provides a uniform procedure for users to choose their own notation.

A semantic network must have an *algorithm for encoding information*. This is provided for by the interfaces to SNePS, e.g., the parser component of our ATN parser-generator inputs English sentences and outputs SNePS networks.

A semantic network must have an *"assimilation"* mechanism for building new information in terms of stored information. SNePS provides for this by the Uniqueness Principle, which enforces node sharing during network building. The assimilation is demonstrated by the generator component of our ATN parser-

generator, which takes SNePS nodes as input and produces English output expressing those nodes. In our conversation with CASSIE, the node built to represent the new fact, 'Lucy is sweet', was expressed in terms of the already existing node for Lucy (who had previously been described as young) by 'young Lucy is sweet'.

A semantic network should be *neutral* with respect to network formalisms at higher levels in the Brachman hierarchy. SNePS is a semantic network at the "logical" level, and CASSIE is (perhaps) at the "conceptual" level. SNePS is neutral in the relevant sense; it is not so clear whether CASSIE is. But a more important issue than neutrality is the reasons why one formalism should be chosen over another. Several possible criteria that a researcher might consider are: *efficiency* (including the ease of interfacing with other modules; e.g., our ATN parser-generator has been designed for direct interfacing with SNePS), *psychological adequacy* (irrelevant for SNePS, but precisely what CASSIE is being designed for), *ontological adequacy* (discussed below), *logical adequacy* (guaranteed for SNePS, because of its inference package), and *natural-language adequacy* (a feature of SNePS's interface with the ATN grammar).

A semantic network should be *adequate* for any higher-level network formalism. SNePS meets this nicely: KL-ONE can be implemented in SNePS (Tranchell 1982).

A semantic network should have a *semantics*. We presented that in Sect. 3.3. But there are at least two sorts of semantics. SNePS nodes have a meaning *within the system* in terms of their links to other nodes; they have a meaning *for users* as provided by nodes at the heads of LEX arcs. Arcs, on the other hand, only have meaning within the system, provided by node- and path-based inference rules (which can be thought of as procedures that operate on the arcs). In both cases, there is an "internal", system's semantics that is holistic and structural: the meaning of the nodes and arcs are not given in isolation, but in terms of the entire network. This sort of "syntactic" semantics differs from a semantics that provides links to an external interpreting system, such as a user or the "world"—i.e., links between the network's way of representing information and the user's way. It is the latter sort of semantics that we provided for CASSIE with respect to an ontology of Meinongian objects, which are not to be taken as representing things in the world. CASSIE's ontology is an *epistemological ontology* (Rapaport 1985/1986) of the purely intensional items that enable a cognitive agent to have beliefs about the world. It is a theory of what there must be in order for a cognitive agent to have beliefs about what there is.

REFERENCES

(1)   Brachman, R. J. (1977), "What's in a concept: structural foundations for semantic networks," *Int. J. Man-Machine Studies* 9: 127-52.

(2)   Brachman, R. J. (1979), "On the Epistemological Status of Semantic Networks," in N. V. Findler (ed.), *Associative Networks* (New York: Academic Press): 3-50; reprinted in (3): 191-215.

(3)   Brachman, R. J., & Levesque, H. J. (1985), *Readings in Knowledge Representation* (Los Altos, CA: Morgan Kaufmann).

(4)   Brachman, R. J.; Fikes, R. E.; & Levesque, H. J. (1983), "KRYPTON: Integrating Terminology and Assertion," *Proc. AAAI-83*: 31-35.

(5)   Brachman, R. J.; Gilbert, V. P.; & Levesque, H. J. (1985), "An Essential Hybrid Reasoning System: Knowledge and Symbol Level Accounts of KRYPTON," *Proc. IJCAI-85*: 532-39.

(6)   Carnap, R. (1928), *The Logical Structure of the World*, R. A. George (trans.) (Berkeley: U. California Press, 1967).

(7)   Castañeda, H.-N. (1974), "Thinking and the Structure of the World," *Philosophia* 4: 3-40.

(8)   Castañeda, H. N. (1975), *Thinking and Doing* (Dordrecht: D. Reidel).

(9)   Fine, K. (1983), "A Defence of Arbitrary Objects," *Proc. Aristotelian Soc.*, Supp. Vol. 58: 55-77.

(10)   Lambert, K. (1983), *Meinong and the Principle of Independence* (Cambridge, Eng.: Cambridge U. Press).

(11)   Maida, A. S., & Shapiro, S. C. (1982) "Intensional Concepts in Propositional Semantic Networks," *Cognitive Science* 6: 291-330; reprinted in (3): 169-89.

(12)   McCarthy, J. (1979), "First Order Theories of Individual Concepts and Propositions," in J. E. Hayes, D. Michie, and L. Mikulich (eds.), *Machine Intelligence* 9(Chichester, Eng.: Ellis Horwood): 129-47; reprinted in (3): 523-33.

(13)   McKay, D. P., & Shapiro, S. C. (1981), "Using Active Connection Graphs for Reasoning with Recursive Rules," *Proc. IJCAI-81*: 368-74.

(14)   Meinong, A. (1904), "Über Gegenstandstheorie," in R. Haller (ed.), *Alexius Meinong Gesamtausgabe*, Vol. II (Graz, Austria: Akademische Druck- u. Verlagsanstalt, 1971): 481-535. Eng. trans. ("The Theory of Objects") by I. Levi et al., in R. M. Chisholm (ed.), *Realism and the Background of Phenomenology* (New York: Free Press, 1960): 76-117.

(15)   Parsons, T. (1980), *Nonexistent Objects* (New Haven: Yale U. Press).

(16)   Rapaport, W. J. (1978), "Meinongian Theories and a Russellian Paradox," *Noûs* 12: 153-80.

(17)   Rapaport, W. J. (1981), "How to Make the World Fit Our Language: An Essay in Meinongian Semantics," *Grazer Phil. Studien* 14: 1-21.

(18)   Rapaport, W. J. (1984b), "Belief Representation and Quasi-Indicators," Tech. Rep. 215 (SUNY Buffalo Dept. of Computer Science).

(19)   Rapaport, W. J. (1985a), "Meinongian Semantics for Propositional Semantic Networks," *Proc. 23rd Annual Meeting Assoc. Computational Linguistics* (U. Chicago): 43-48.

(20)   Rapaport, W. J. (1985/1986), "Non-Existent Objects and Epistemological Ontology," *Grazer Phil. Studien*, Vol. 25/26.

(21)   Rapaport, W. J., & Shapiro, S. C. (1984), "Quasi-Indexical Reference in Propositional Semantic Networks," *Proc. 10th Int. Conf. Comp. Ling.* (COLING-84) (Stanford U.): 65-70.

(22)   Routley, R. (1979), *Exploring Meinong's Jungle and Beyond* (Canberra: Australian Nat'l. U., Dept. of Philosophy).

(23)   Shapiro, S. C. (1978), "Path-Based and Node-Based Inference in Semantic Networks," in D. Waltz (ed.), *Tinlap-2: Theoretical Issues in Natural Language Processing* (New York: ACM): 219-25.

(24)   Shapiro, S. C. (1979), "The SNePS Semantic Network Processing System," in N. V. Findler (ed.), *Associative Networks* (New York: Academic Press): 179-203.

(25)   Shapiro, S. C. (1982), "Generalized Augmented Transition Network Grammars for Generation from Semantic Networks," *American J. Computational Linguistics* 8: 12-25.

(26)   Shapiro, S. C., & McKay, D. P. (1980), "Inference with Recursive Rules," *Proc. AAAI-80*: 151-53.

(27)   Shapiro, S. C., & Rapaport, W. J. (1985), "SNePS Considered as a Fully Intensional Semantic Network," in G. McCalla & N. Cercone (eds.), *Knowledge Representation* (Berlin: Springer-Verlag, forthcoming); also, Tech. Rep. 85-15 (SUNY Buffalo Dept. of Computer Science, 1985).

(28)   Shapiro, S. C.; Martins, J.; & McKay, D. P. (1982), "Bi-Directional Inference," *Proc. 4th Annual Conf. Cognitive Science Soc.* (U. Michigan): 90-93.

(29)   Shapiro, S. C.; Srihari, S. N.; Geller, J.; & Taie, M.-R. (1986), "A Fault Diagnosis System Based on an Integrated Knowledge Base," *IEEE Software* 3.2(March 1986)48-49.

(30)   Srihari, R. K. (1981), "Combining Path-Based and Node-Based Inference in SNePS," Tech. Rep. 183 (SUNY Buffalo Dept. of Computer Science).

(31)   Tranchell, L. M. (1982), "A SNePS Implementation of KL-ONE," Tech. Rep. 198 (SUNY Buffalo Dept. of Computer Science).

(32)   Woods, W. A. (1975), "What's in a Link: Foundations for Semantic Networks," in D. G. Bobrow & A. M. Collins (eds.), *Representation and Understanding* (New York: Academic Press): 35-82; reprinted in (3): 217-41.

(33)   Zalta, E. N. (1983), *Abstract Objects* (Dordrecht: D. Reidel).