DEIXIS IN NARRATIVE
A Cognitive Science Perspective

Edited by

Judith F. Duchan
Gail A. Bruder
Lynne E. Hewitt
State University of New York at Buffalo

4

AN INTRODUCTION TO A
COMPUTATIONAL READER
OF NARRATIVES

Stuart C. Shapiro
William J. Rapaport
State University of New York at Buffalo

In this chapter, we describe the SNePS knowledge-representation and reasoning system. We look at how SNePS is used for cognitive modeling and natural language competence. SNePS has proven particularly useful in our investigations of narrative understanding. Several other chapters in this book (Almeida; Rapaport & Shapiro; Yuhan & Shapiro) use SNePS to discuss specific issues in areas relevant to narrative understanding.

SNePS

SNePS is an intensional, propositional, semantic-network knowledge-representation and reasoning system that is used for research in artificial intelligence (AI) and in cognitive science. "Knowledge" representation is the study of the representation of information in an AI system (because the information need not be true, a more accurate name would be "belief" representation; cf. Rapaport, 1992).

SNePS (Shapiro, 1971, 1979; Shapiro & Rapaport, 1987, 1992) is a programming language whose primary data structure is a semantic network (a labeled, directed graph), with commands for building such networks and finding nodes in such a network given arbitrary descriptions (including partial descriptions). The particular kind of semantic network SNePS builds is *propositional*. Furthermore, the particular kind of propositional semantic network that we believe is appropriate for cognitive modeling and natural-language competence is one that is *fully intensional*. Both of these notions will be clarified later (cf. Maida & Shapiro, 1982; Shapiro & Rapaport, 1987, 1991).

In addition to being able to build and find networks, the SNePS Inference Package (SNIP) permits *node-based reasoning, path-based reasoning,* and *belief revision*. Node-based reasoning can be thought of as conscious reasoning following explicit rules stated in the form of networks. These "rules" are really axioms or nonlogical postulates, not rules of inference. The rules of inference are not explicitly represented, but are implemented in the "inference engine"—SNIP. Path-based reasoning can be thought of as subconsious reasoning; it is a generalization of the notion of inheritance found in many other semantic-network systems (cf. Shapiro, 1978, 1991; Srihari, 1981). A belief revision system is a facility for detecting and removing inconsistent beliefs; the SNePS version is based on a form of relevance logic (cf. Martins & Shapiro, 1988).

Finally, in addition to using the SNePS User Language (a programming language written in Lisp) to build, find, and deduce information directly, one can also interact with SNePS in natural language, using a generalized augmented-transition-network parser-generator (Shapiro, 1982). This makes it especially appropriate for use in our project of understanding narrative text.

PROPOSITIONAL SEMANTIC NETWORKS

A Brief Introduction to Semantic Networks

A semantic network is usually thought of as a labeled, directed graph, whose nodes represent entities and whose arcs represent binary relations between the entities. In the stereotypical semantic network found in most AI texts, there are arcs with such labels as "ISA" and "A-KIND-OF," corresponding, roughly, to set membership and the subset relation. When these labels prove insufficient, labels such as "PROPERTY" are sometimes used, corresponding to the relation of having-as-a-property.

Such networks, which can be thought of as descendents of the medieval Porphyrian tree, are essentially taxonomic and "object-oriented": The nodes represent individuals, classes, or other objects, as well as their properties. One important feature of such networks is the property of *inheritance*: Information stored about an object represented by some node need not be stored (redundantly) at all the nodes that are related to it by ISA or A-KIND-OF. For example, if a dog ISA mammal, and mammal is A-KIND-OF animal, then a dog ISA animal; and if animal has as a PROPERTY mortal, then dog has as a PROPERTY mortal; that is, if a dog is a mammal, mammals are animals, and animals are mortal, then a dog inherits being an animal and being mortal. There are, of course, much more sophisticated varieties of such taxonomic inheritance hierarchies (such as the KL–ONE family of systems), and their logical properties have been extensively investigated (for surveys, see Brachman & Levesque, 1985; Lehmann, 1992; Sowa, 1991, 1992).

Representing Propositions

In a *propositional* semantic network, by contrast, the nodes can represent propositions in addition to objects, classes, and properties. To see how this can be done, consider the network in Fig. 4.1. Suppose that the Plato node represents Plato, the philosopher node represents the class of philosophers, and the isa arc represents the relationship that holds between Plato and the class of philosophers. Now suppose that one wanted to deny that Plato was a philosopher. How would negation be represented? One couldn't, using only graph-theoretical techniques, easily or simply negate the Plato node or the philosopher node. The former option would presumably mean that not-Plato is a philosopher, but what is a "not-Plato"? The latter option would presumably mean that Plato is a not-philosopher, but what is a "not-philosopher"? Nor could one negate the isa arc without violating the conventions of graph theory.

The SNePS solution is to "split" the arc in two. There are several ways this can be done, one of which is shown in Fig. 4.2. The Plato node can still represent Plato, and the philosopher node can still represent the class of philosophers. But now there is a third node, arbitrarily labeled M1, that can represent the *proposition* that Plato is a philosopher. This node is "structured": It represents the proposition as explicitly consisting of an object (Plato), and a class (of philosophers). The class membership relation represented by the arcs is implicit.

Suppose, however, that we wanted to make the relation that holds between Plato and the class of philosophers explicit, so that we could talk about it (in SNePS, one can only talk about nodes, not arcs). There are several ways this could be done; one way is shown in Fig. 4.3. Here, the Plato and philosopher nodes can be interpreted as before, as can node M1. The new ISA node represents the class membership relation itself. The arc labels reflect the new structure of this proposition. In this network, the ISA relation is explicit, but the higher order rel/object-1/object-2 ternary relation is implicit.

Case frames (cf. Fillmore, 1968) such as rel/object-1/object-2 can also be thought of along the lines of Davidson's analysis of events (Davidson, 1967): The proposition that Plato is a philosopher is represented by the network of Fig. 4.3 as being structured as follows:

$$\exists p[\text{rel}(p, \text{ISA}) \ \& \ \text{object-1}(p, \text{Plato}) \ \& \\ \text{object-2}(p, \text{philosopher})],$$

with M1 being thought of as the value of p or (better) as a Skolem constant. (For more details on the nature of the case frames in SNePS, cf. Shapiro & Rapaport,

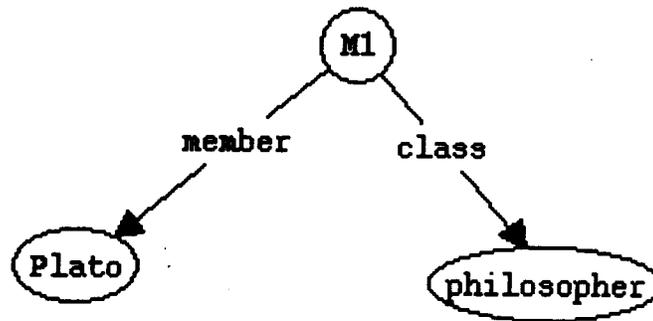FIG. 4.1. An ISA network representation.

FIG. 4.2.   One possible SNePS representation.

1987; Shapiro, 1991. For a view of proposition nodes as functional terms, cf. Shapiro, 1993. In using SNePS for cognitive modeling, we use both the implicit case frame of Fig. 4.2 and the explicit case frame of Fig. 4.3, as discussed later.)

Advantages of Propositional Semantic Networks

Representing Beliefs. To see the advantage of propositional semantic networks over object-oriented ones, consider how we would represent the proposition:

Mary believes that John is rich, but he isn't.

One way of representing this in SNePS is shown in Fig. 4.4. Nodes that are marked with an exclamation point are said to be *asserted*; they represent "beliefs" of the system. Node M2 ! represents the system's belief that an agent, Mary, performs the mental act of believing, directed (in a Meinongian sense) to the object M1 (cf. Rapaport, 1978). Node M1 represents the proposition that John is rich; thus, M2 ! represents the system's belief that Mary believes that John is rich. Node M3 ! represents the system's belief that it is not the case that M1; thus, M3 ! represents the "but he isn't" part of the proposition. (The min/max/arg case frame is used to represent the proposition that at least min and at most max of the propositions pointed to by argument arcs are true; here,
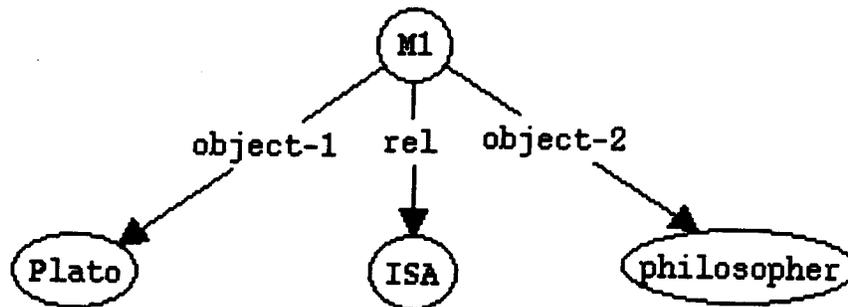

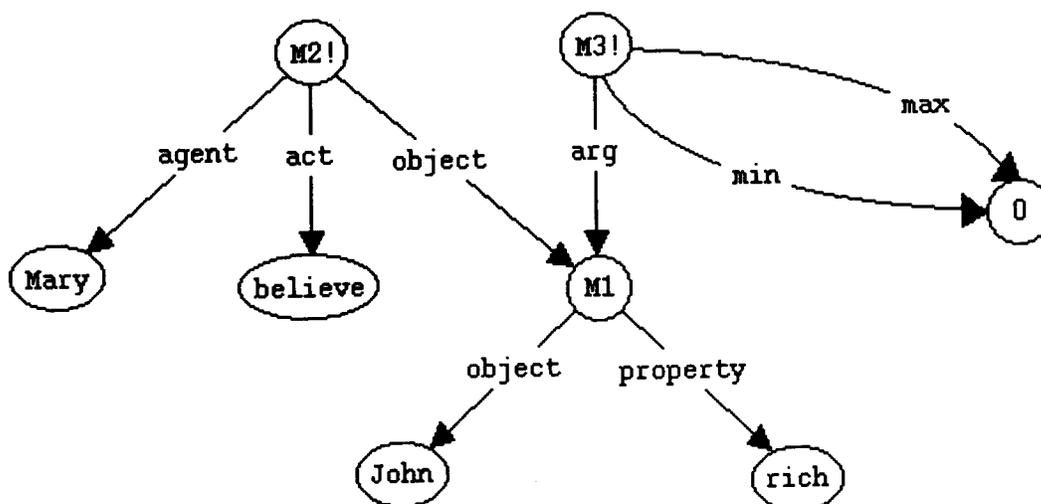FIG. 4.3.   Another possible SNePS representation.

FIG. 4.4. One SNePS representation of: Mary believes that John is rich, but he isn't.

between 0 and 0 of the single argument M1 are true—that is, M1 is false. For details, cf. Shapiro, 1979; Shapiro & Rapaport, 1987.)

For ease of exposition, we will let [[n]] represent the denotation of node n—what n represents. For example, referring again to Fig. 4.4, we could say that [[M2!]] is the proposition that [[Mary]] believes the proposition [[M1]], and [[M1]] is the proposition that [[John]] has the property [[rich]].

Actually, the network of Fig. 4.4 is vastly oversimplified. It fails to take into account the difference between *de re* and *de dicto* belief reports, for one thing. A more accurate representation of Mary's mistaken belief that John is rich is shown in Fig. 4.5. Node M9! represents the system's belief that an agent named "Mary" believes that something (represented by B2) is rich. In particular, [[M2!]] is the proposition that [[B1]] has the propername [[M1]], expressed in English by the lexical item "Mary." Node M8 represents the object of Mary's belief, namely, that [[B2]] is an object with the property [[M7]], which is expressed in English by "rich." Node M6! represents the system's belief that Mary believes that [[B2]]—the rich object—is named "John." So, taken together, M6! and M9! represent the system's belief that Mary believes *de dicto* that John is rich. Finally, M10! represents the system's belief that it is not the case that John is rich. Consider the agent/act/object case frames whose act is [[M3]]. The subnetworks at the heads of the object arcs form the "belief space" of the agent. They represent the system's beliefs about the agent's beliefs.[1]

---

[1]Details of the representation of belief reports and belief spaces is given in Rapaport, 1986a; Rapaport, Shapiro, & Wiebe, 1986; and Wiebe & Rapaport, 1986. A discussion of what the nodes in such belief spaces represent can be found in Shapiro, 1993, and Shapiro & Rapaport, 1991. Recent papers by Crimmins and Perry seem to be consistent with the ontology behind SNePS networks. Their *beliefs, ideas,* and *notions* seem to correspond precisely to SNePS propositional nodes, property nodes, and individual nodes, respectively (cf. Crimmins, 1989; Crimmins & Perry, 1989).
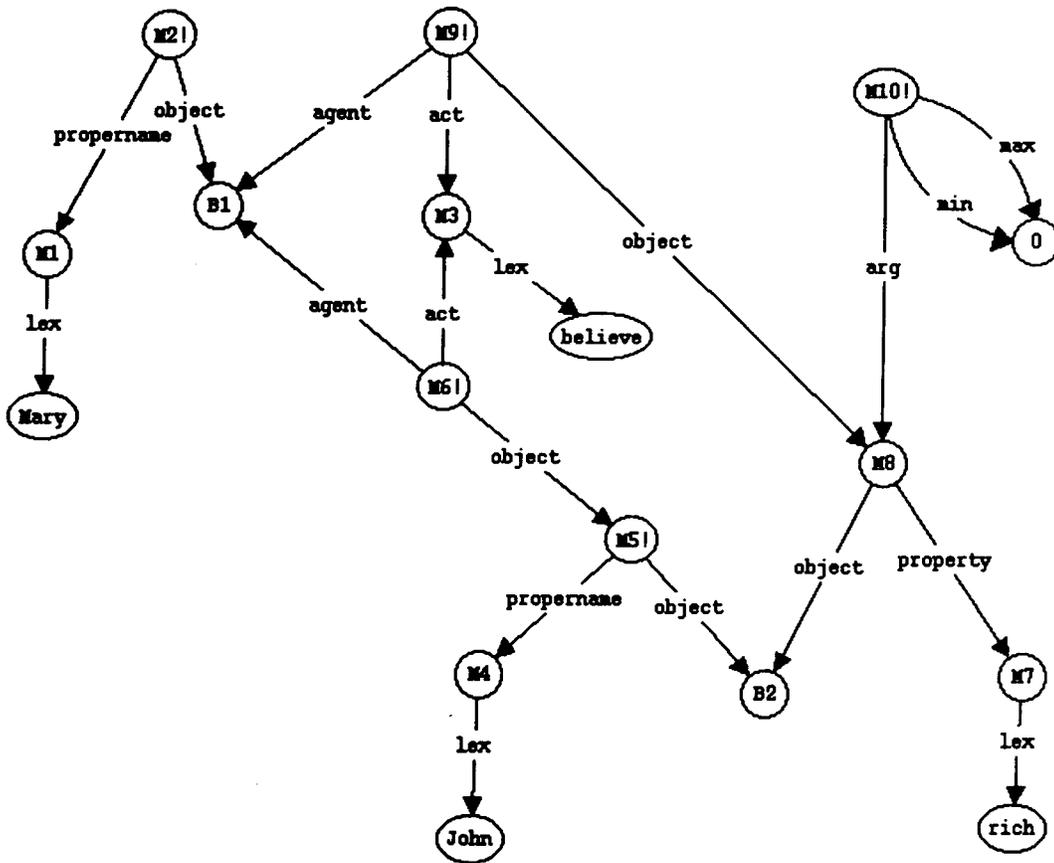
FIG. 4.5. A more accurate SNePS representation of: Mary believes that John is rich, but he isn't.

Node-Based Reasoning. Another advantage of propositional semantic networks is their use in node-based reasoning. Rules can be represented as propositions in the same knowledge base as the propositions to which the rules are intended to apply. This differs from the architecture of the typical production system used to implement an expert system. In such architectures, the rules are in long-term memory separate from the working memory containing the propositions the rules manipulate.

SNIP interprets certain propositional nodes as being rules and performs forward- and backward-inference using them. For instance, the proposition:

$$\forall v1[human(v1) \rightarrow mortal(v1)]$$

might be represented by node M1 ! in Fig. 4.6. It has a universal quantifier arc (forall) pointing to a variable node, V1; an antecedent arc pointing to node P1; and a consequent arc pointing to node P2. Nodes P1 and P2 can be thought of as propositional functions in Russell's sense; they are called *pattern* nodes in SNePS. Node P1 represents the propositional function that [[V1]] is human; node P2 represents the propositional function that [[V1]] is mortal.
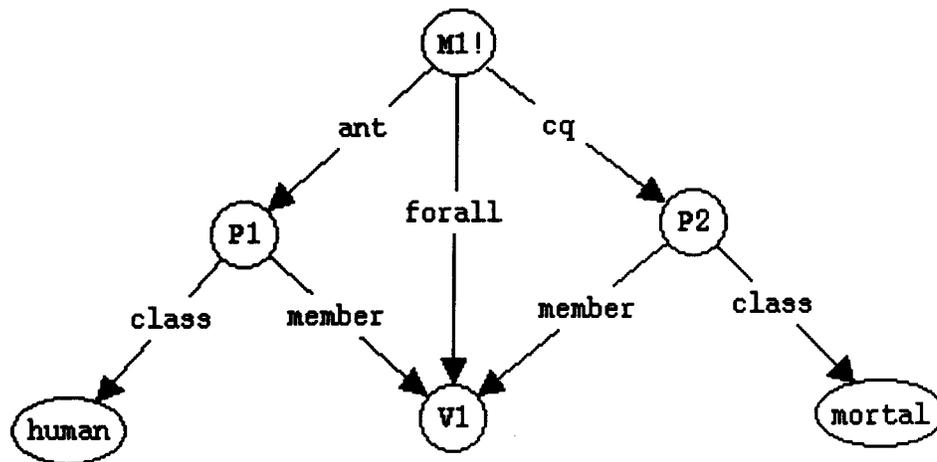
FIG. 4.6. A SNePS representation of $\forall v1[\text{human}(v1) \rightarrow \text{mortal}(v1)]$.

Suppose that the system is told that Socrates is human, resulting in node M2 ! being asserted (see Fig. 4.7). Note that the human node is shared by nodes M2 ! and P1; this is a result of the *uniqueness principle*: Every node represents some entity in the domain of discourse, and no two nodes represent the same entity (cf. Maida & Shapiro, 1982; Shapiro & Rapaport, 1987). Next, suppose that the system is asked who is mortal. If the question is phrased in such a way that the system is merely being asked to *find* asserted nodes representing propositions of the form "*x* is mortal," it will find none. However, if the question is phrased in such a way that the system is asked to *deduce* whether anyone is mortal, it will behave as follows: Imagine that it has been asked to find or to build asserted propositional nodes matching (i.e., with the structure of) the pattern node P2.
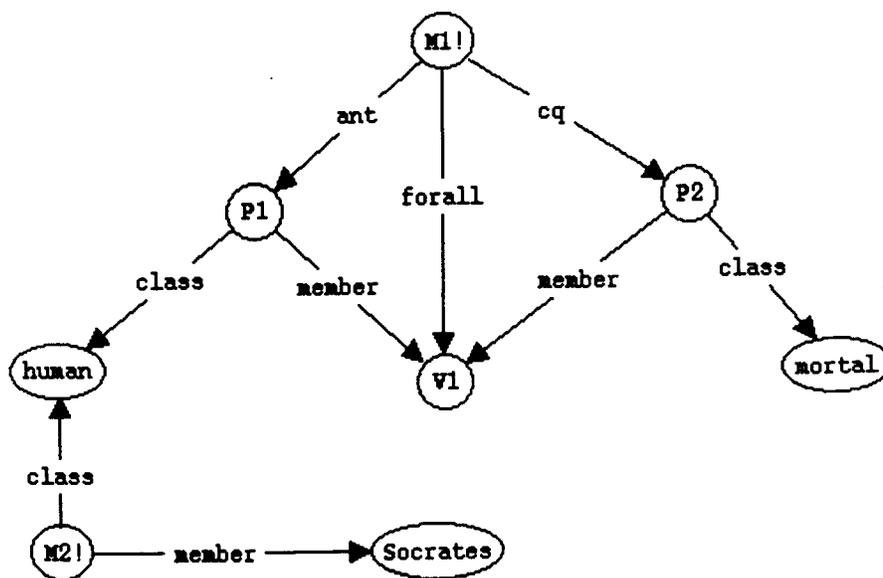

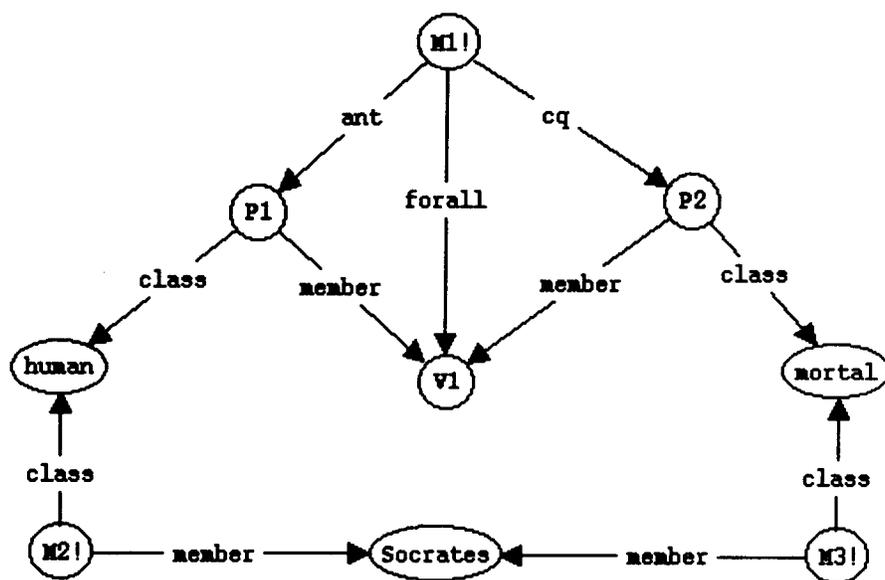FIG. 4.7. After telling the system that Socrates is human.

FIG. 4.8.   After inferring that Socrates is mortal.

Finding none (because there aren't any), it will backchain and seek asserted propositional nodes matching P1. It will find M2!, which matches P1 if the Socrates node is bound to V1. This information is used to assert a new node, M3!, matching P2, with V1 bound to the Socrates node. Node M3! is the conclusion of the inference. The network has now grown, to look like Fig. 4.8.

INTENSIONAL KNOWLEDGE REPRESENTATION

SNePS is an intensional knowledge-representation system; that is, it supports multiple representations of what could be one physical object.[2] We argued in earlier papers that the nodes of a semantic network not only can, but ought to, represent intensional entities (Maida & Shapiro, 1982; Rapaport, 1985a, in press; Shapiro & Rapaport, 1987, 1991). By *intensional entity*, we have in mind things like Fregean senses (Frege, 1892), Meinongian objects (Meinong, 1904), Castañedian guises (Castañeda, 1972), and Routleyan items (Routley, 1979). There does not seem to be a clear characterization in the literature of what these things are, but they seem to satisfy the following:

1. They are nonsubstitutable in intensional contexts, even if they are "the same" (i.e., they can be equivalent without being identical). The morning star and the evening star are examples.

---

[2]We owe this way of putting the matter to Susan Haller.

2. They can be indeterminate. That is, they can be "incomplete", as in the case of fictional entities. (Did Sherlock Holmes have a mole on his left arm or not?)

3. They need not exist. The nonexisting golden mountain is an example.

4. They need not be possible. The round square is an example.

5. They can be distinguished even if necessarily identical. For example, *the sum of 2 and 2* and *the sum of 3 and 1* are distinct objects of thought.

We claim that to model a mind, a knowledge-representation and reasoning system must model only intensional entities. There are two main arguments for this. The first may be called *The Argument From Fine-Grained Representation* and is summarized as follows: Inten*T*ional entities (i.e., objects of thought) are inten-*S*ional. That is, one can have two objects of thought that correspond to only one extensional object (as in the familiar examples of the morning star and the evening star). The second argument may be called *The Argument From Displacement*, summarized thus: We can think and talk about nonexistent objects—fictional ones, impossible ones, etc.—thus, we need to be able to represent and reason about them, especially if we are interested in using the system for understanding works of fiction.

As an example of how intensional objects can be represented in SNePS, consider Fig. 4.9. It shows one of the possible networks that can represent "The Morning Star has the property of being a planet." Node M2! represents the proposition that intensional entity [[B1]] is the Morning Star. The proposition is structured as follows: [[B1]] is an object whose propername is [[M1]], and [[M1]] is expressed in English by the lexical item "The Morning Star." (This is neither the only or even the best way to represent this proposition, but it will suffice.) Node M4! represents the proposition that [[B1]] is a planet. This proposition is structured as follows: [[B1]] is an object that has the property
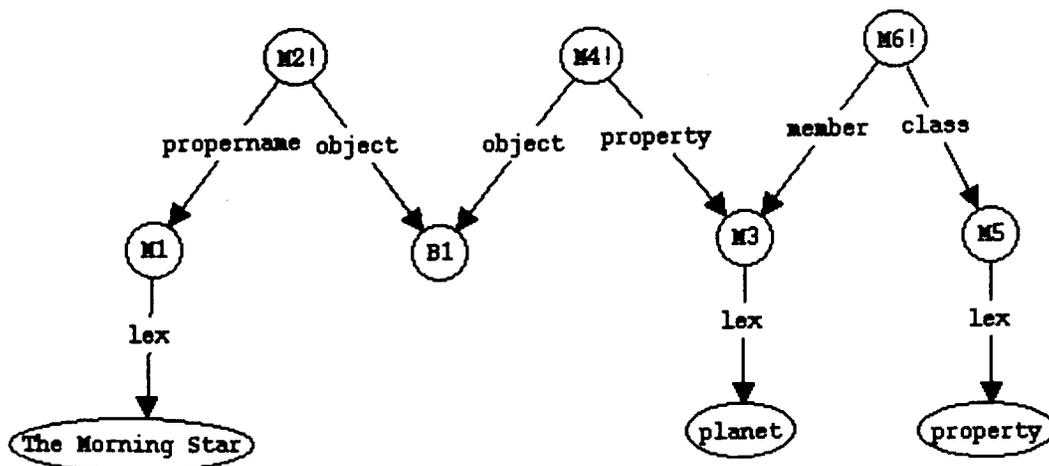

FIG. 4.9. A possible SNePS representation of: The Morning Star has the property of being a planet.

[[M3]], which is expressed in English by the lexical item "planet." The network described thus far represents "The Morning Star is a planet." To make it explicit that being a planet is a property, we use node M6!, whose structure is that [[M3]] is a member of the class [[M5]], expressed in English by "property."

Note that *being the Morning Star* is predicated of [[B1]], and that [[B1]] can be considered as a bare particular (cf. Allaire, 1963, 1965; Landman, 1986). Similarly, *being a planet* is predicated of [[B1]]. Finally, *being a property* is predicated of [[M3]] (though [[M3]] is not quite bare; it has structure, as do [[M1]] and [[M5]], all of which are structured objects).

If the system is next told, "The Morning Star is the Evening Star," the structure shown in Fig. 4.10 will be added. Here, [[B2]] is the intensional entity *The Evening Star*, which is distinct from [[B1]]. (Proposition [[M8!]] asserts that [[B2]] is named "The Evening Star.") Node M9! represents the proposition that [[B1]] and [[B2]] are equivalent, that is, the Morning Star is the Evening Star. There is no node that represents the extensional entity that is both the Morning and Evening Stars. (You might think that it would be a node representing an entity named "Venus," but that would just be a third intensional entity, equivalent to the other two. For a very different way of representing intensionality in SNePS, see Wyatt, 1989, 1990, 1993.)

PUTTING IT ALL TOGETHER: CASSIE READS A NARRATIVE

We have been talking about the system as if it were (merely) an AI program. It is that, of course, but it is, in particular, a program in the area of AI called *cognitive modeling*. We are constructing a model of a cognitive agent who can reason, solve
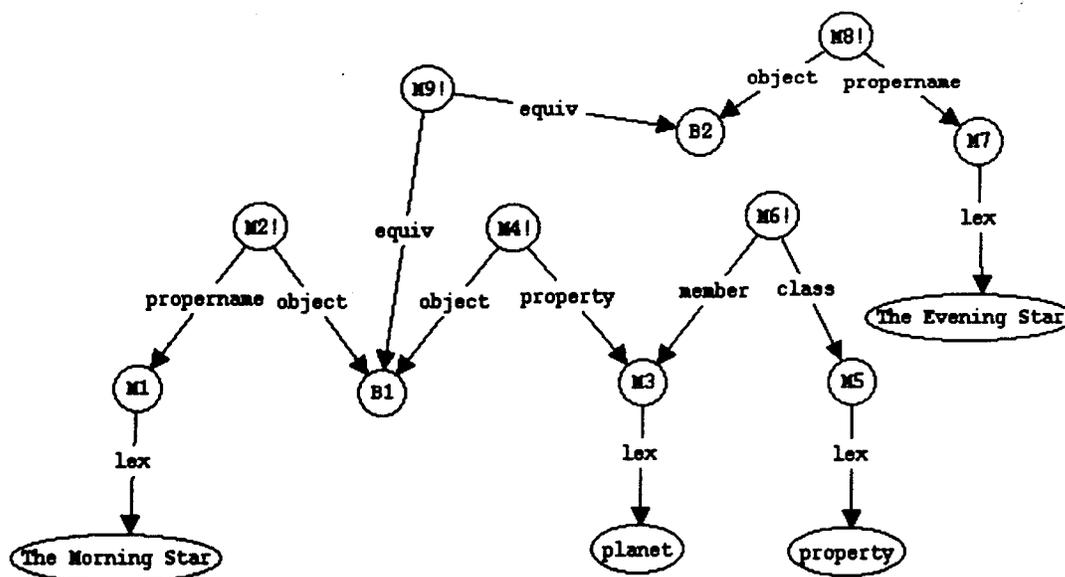

FIG. 4.10. After telling the system that: The Morning Star is the Evening Star.
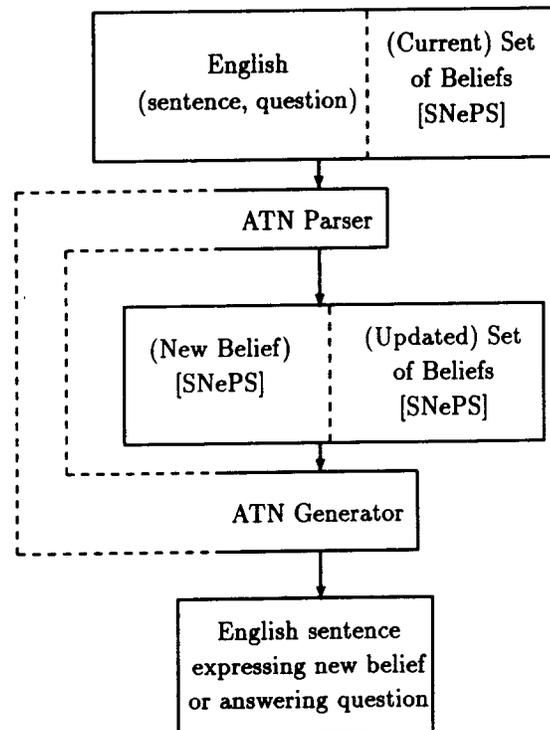
FIG. 4.11. Communicating with Cassie.

problems, and read and converse in natural language. We argued elsewhere that such a model of a cognitive agent *is* a cognitive agent (or, more modestly, that such a model will be such an agent when all of the problems of cognitive modeling will have been solved; cf. Rapaport, 1986b, 1988a, 1988b; Shapiro & Rapaport, 1991).

We call our cognitive agent *Cassie* (the *C*ognitive *A*gent of the *S*NePS *S*ystem—an *I*ntelligent *E*ntity). SNePS nodes represent the objects of Cassie's thoughts. In addition, we can give a fully intensional semantic interpretation for the nodes in terms of a Meinongian Theory of Objects without recourse to possible worlds—which are neither needed or wanted, because we want to be able to deal with impossible objects (cf. Rapaport, 1985a, in press; Shapiro & Rapaport, 1987).

Cassie can read a narrative in (a fragment of) English, as indicated in Fig. 4.11. Cassie has a set of background or current beliefs, represented in SNePS, and reads or is asked something in English. The English statement or question is interpreted by the parser fragment of a generalized augmented-transition-network parser-generator (Shapiro, 1982), which outputs a SNePS node. This node might be an already-existing node in Cassie's "mind" (cf. Rapaport, 1988a) or a new node asserted as the result of an inference (as described previously). The result is a possibly updated set of beliefs in Cassie's mind (her mind may, literally, have changed). The node that was found or asserted is then expressed in English by the generator fragment of the generalized augmented-transition-network parser-generator. Figure 4.12 shows Cassie reading a short narrative, and occasionally being asked questions to test her understanding. The sentence on the first line of each interaction, following the ":" prompt, is either a sentence of the narrative or a question posed to Cassie. Cassie's response begins on the subsequent line.

```
1: John petted a yellow cat.
I understand that John petted a yellow cat.
Time (sec.): 0.617

2: The cat is a manx.
I understand that the yellow cat is a manx.
Time (sec.): 0.25

3: What did John pet?
John petted a yellow cat.
Time (sec.): 0.7

4: The dog is a mammal.
I understand that dogs are mammals.
Time (sec.): 0.283

5: Every dog is owned by a boy.
I understand that every dog is owned by a boy.
Time (sec.): 0.35

6: Young Lucy petted a yellow dog.
I understand that young Lucy petted a yellow dog.
Time (sec.): 1.017

7: The boy saw her.
I understand that the boy saw young Lucy.
Time (sec.): 0.6
```

FIG. 4.12.   A conversation with Cassie.

Before she began reading, Cassie's mind contained one explicit belief (node M1! in Fig. 4.13), some implicit beliefs for path-based reasoning, and some other, deictic, information. The implicit beliefs allow certain kinds of inheritance, for instance, that a path consisting of a SUPERCLASS arc followed by zero or more pairs of converse-SUBCLASS/SUPERCLASS arcs can itself be considered as a ("virtual") SUPERCLASS arc; no more will be said of them here. The deictic information consists of a "now"-pointer pointing to node B1, representing the current time, and an "I"-pointer pointing to node B2, representing Cassie's "self-concept." The now-pointer–mechanism is a fairly primitive way of representing time and tense; it was superseded by the more sophisticated facility described in Almeida (1987). The I-pointer–mechanism was described more fully in Rapaport, Shapiro, and Wiebe (1986).

Node M1! represents Cassie's belief that:

$$(\forall v1,v2,v3)[\text{if } v1 \text{ is a } v2 \text{ \& } v2 \text{ is a } v3, \text{ then } v1 \text{ is a } v3].$$

But note that this is represented, not as the transitivity of "is a", but using three different propositions. Node P1 represents the proposition that V1 is a V2 in the sense of class membership; node P2 represents the proposition that V2 is a V3
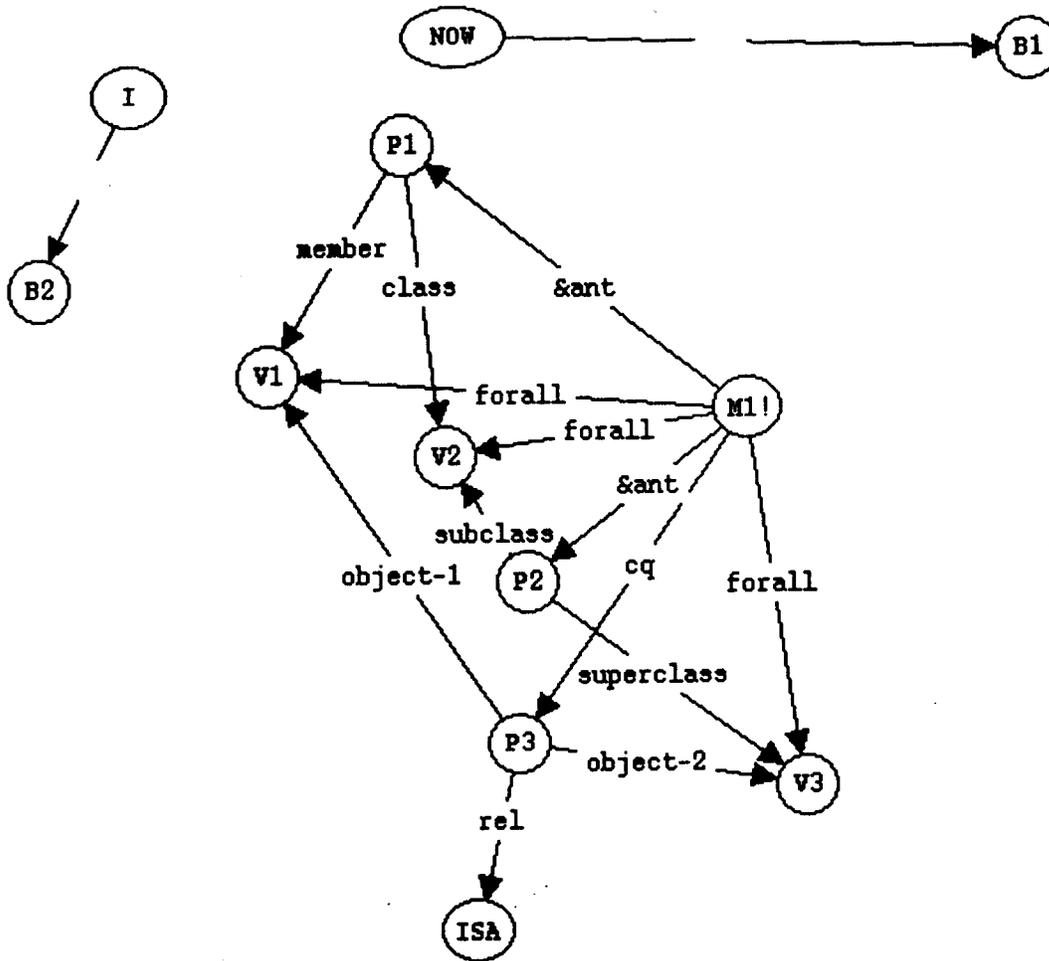
FIG. 4.13. Cassie's mind before the conversation begins.

in the sense of the SUBCLASS/SUPERCLASS relation; and node P3 represents the proposition that V1 is a V3 using the ISA relation of taxonomic semantic networks. Why the variety? The choices for P1 and P2 should be apparent. Class membership is not used for P3 for the following reason: We use class membership to represent the relation obtaining between an individual and the basic-level category to which it belongs (in the sense of Rosch, 1978). We use the SUB-CLASS/SUPERCLASS relation to represent the relation obtaining between a basic-level category and its superordinate-level category. And we use the ISA relation to represent the relation obtaining between an individual and any non-basic-level categories (subordinate or superordinate). One advantage of this for the purposes of natural-language competence is the ability to get Cassie to converse more normally. (For a more complete explanation of this representation, see Peters & Shapiro, 1987a, 1987b; Peters, Shapiro, & Rapaport, 1988.)

When Cassie reads that John petted a yellow cat, nodes M3 !, M7 !, M8 !, M9 !, M10 !, and M11 ! are asserted, changing her mind to look like Fig. 4.14. Node M3 ! "says" that [[B3]] is named "John"; node M7 ! says that [[B4]] is a cat (*cat*

FIG. 4.14. [[M9 !]] = John petted a yellow cat.

is a basic-level category); node M8! says that [[B4]] (the cat) is yellow; and node M9! says that [[B3]] petted [[B4]]. The past tense "petted" is represented by indicating that the ACT of petting had a StartTIME represented by B5 and an EndTIME represented by B6, that [[B5]] is before [[B6]] (this is asserted by M10!), and that [[B6]] is before "now" (asserted by M11!).

Cassie then reads that *the* cat is a manx. Which cat? Since she only has beliefs about one cat (the one petted by John), she decides that it is the cat that is a manx. Her mind grows to include node M13!, which represents that the cat ([[B4]]) ISA manx (*manx* is a subordinate-level category). Note, too, that Cassie's way of expressing [[M13!]] in English shows us that she understands that the cat that is a manx is the yellow cat that John petted. When we ask Cassie (in order to test her understanding of the narrative that she is reading) what John petted, she responds that he petted a yellow cat rather than that he petted a yellow manx. This might not seem remarkable, but had we not distinguished between the propositions—*being a cat* is represented by class membership, *being a manx* is represented by ISA—Cassie might have replied (and, in an earlier implementation, did reply) that John petted a yellow manx, a decidedly unidiomatic way of putting things in this context. Cassie's mind now looks like Fig. 4.15.

Next, Cassie reads that the dog is a mammal. Which dog? Since there have been none in the narrative thus far, Cassie assumes that this is to be interpreted as a generic sentence, that is, as "dogs are mammals," and node M16! is built. Note that at this point, there are two disconnected subnetworks: That dogs are mammals is entirely irrelevant to Cassie's beliefs about John, the cat, manxes, petting, and being yellow. Well, almost entirely: Cassie does believe that if an individual is a member of a basic-level category that is a subclass of a superordinate-level category, then the individual stands in the ISA relation to the superordinate-level category; because *dog* is a basic-level category that is a subclass of the superordinate-level category *mammal*, there is an implicit connection between the two subnetworks. Cassie's mind now looks like Fig. 4.16.

Cassie then reads that every dog is owned by a boy, and her mental model of the conversation changes to reflect this new piece of information. Node M19!, which represents this rule, has a structure similar to the rule:

$$\forall v8[dog(v8) \rightarrow \exists y[boy(y) \ \& \ owns(y, v8)]].$$

However, the current implementation of SNePS does not have an existential quantifier, so, instead, we use the Skolemized version:

$$\forall v8[dog(v8) \rightarrow [boy(b7(v8)) \wedge owns(b7(v8), v8)]]$$

where b7 is the Skolem function. Specifically, the structure of M19! is: There is a universally quantified variable-node, V8. The ANTecedent of M19! is node P10, which says that [[V8]] is a dog (thus linking up to the previously constructed, isolated subnetwork dominated by node M16!). The ConseQuents of M19! are nodes P12 and P13, both of which refer to P11, which is the result of applying
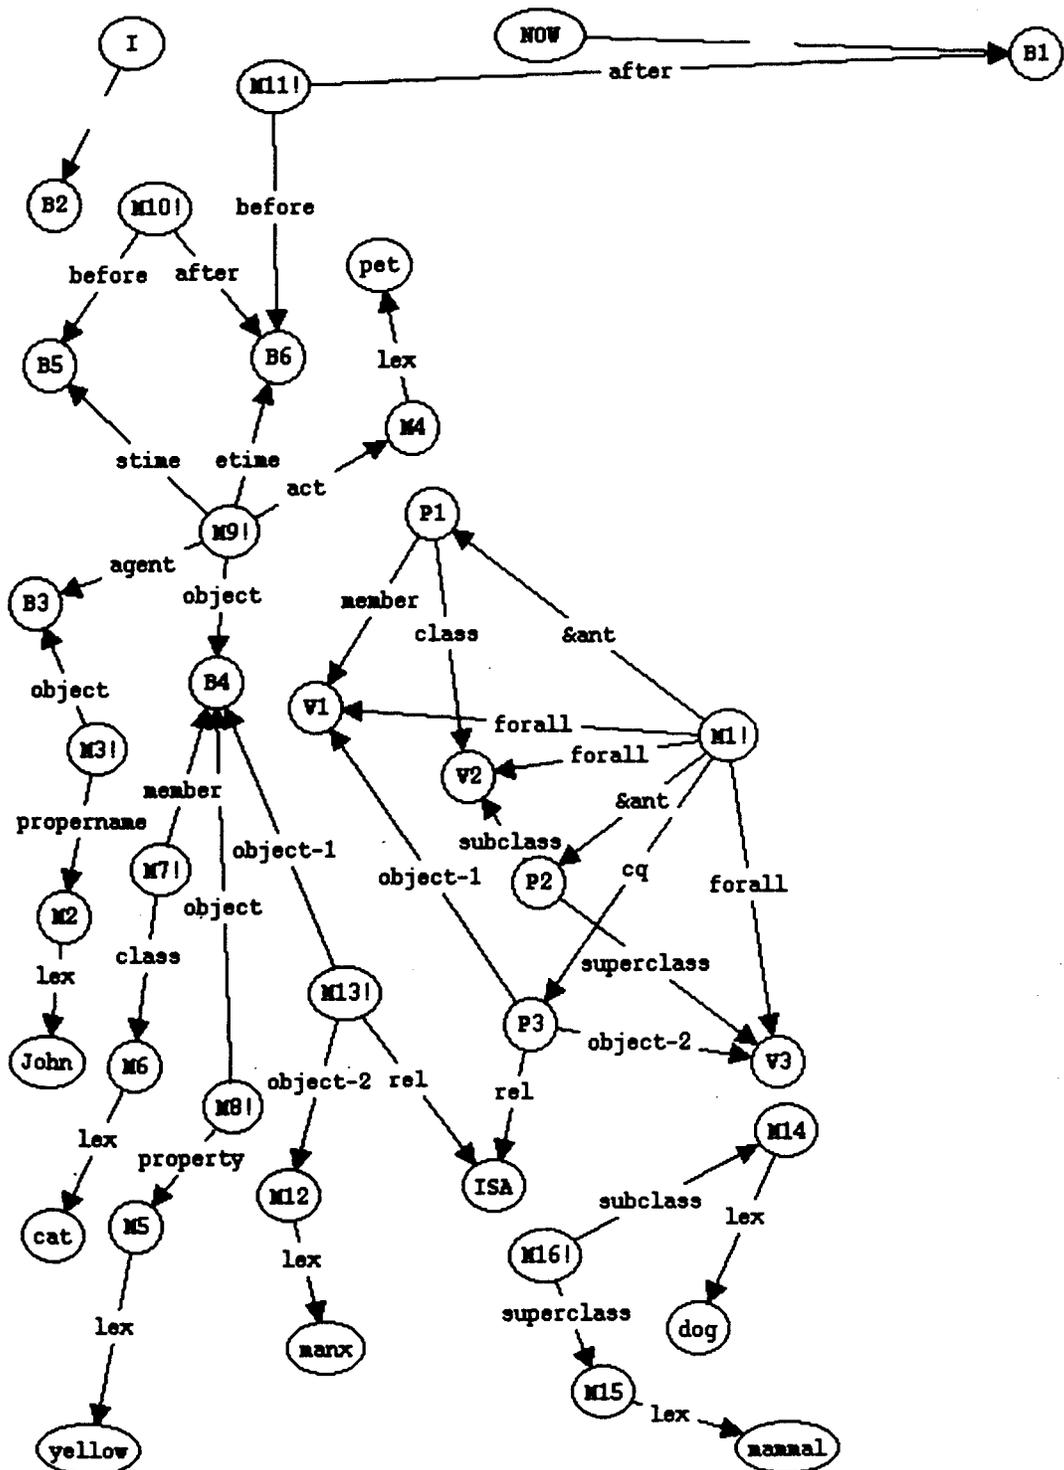
FIG. 4.15. [[M13!]] = The cat is a manx.

FIG. 4.16. [[M16 !]] = Dogs are mammals.

95

the Skolem function B7 to V8. Node P12 says that [[P11]] is a boy; node P13 says that he owns [[V8]]. The present tense of "owns" is represented by having the StartTIME, B8, be before "now", and the EndTIME, B9, be after "now." Cassie's mind now looks like Fig. 4.17.

Next, Cassie reads that young Lucy petted a yellow dog. The first thing that Cassie understands is that someone ([[B10]]) is named "Lucy" and that Lucy is young (nodes M24 ! and M25 !). Then she understands that there is a dog, [[B11]]. So, [[B10]] is young Lucy, and [[B11]] is a dog. Cassie's mind now looks like Fig. 4.18.

However, Cassie already believes that every dog is owned by a boy. So she conceives of an individual, [[M27]], which results from instantiating P11 by substituting B11 for V8, and infers the proposition that this individual is a boy ([[M28 !]]). (Cassie's mind now looks like Fig. 4.19.) Then, Cassie infers that this boy owns the dog ([[M29 !]]). Cassie's mind now looks like Fig. 4.20.

There are other rules that Cassie believes, namely, that dogs are mammals and that if something is a dog, then it is a mammal (rule [[M1 !]]); so she infers that [[B11]] ISA mammal ([[M31 !]]). Cassie's mind now looks like Fig. 4.21.

All of this occurs while she is in the process of understanding that young Lucy petted a yellow dog. What she understands of this sentence so far is that someone is named "Lucy," that that someone is young, and that something else is a dog. She has then inferred that the dog is a mammal and that it is owned by a boy. Now she comes to understand that the dog is yellow ([[M32 !]]). Cassie's mind now looks like Fig. 4.22. Subsequently, she comes to understand that young Lucy petted the yellow dog ([[M33 !]]), and that this event occurred after "now" ([[M35 !]]). Now all previously disconnected subnetworks are linked, and Cassie's mind has grown to look like Fig. 4.23.

Finally, Cassie reads that *the* boy saw her. Which boy? Well, no boys were explicitly talked about, but she has inferred the existence of a boy, so he must be the boy who saw her. Her? Who? An anaphoric-pronoun–resolution system determines that it was young Lucy whom the boy saw (Li, 1986). The result is the set of nodes M37 !, M38 !, and M39 !. Cassie's mind now looks like Fig. 4.24.

The narrative is now complete (actually, this is a fragment of a much longer narrative, but by now you should see how things work). Cassie's mind at the end of the narrative is shown in Fig. 4.24; this is what she came to believe while reading. It is also her mental model of the story of John and Lucy.

SUMMARY

We provided an introduction to the SNePS knowledge-representation and reasoning system, and to Cassie, a computational agent that can read natural language texts and form an understanding of those texts in terms of beliefs represented in
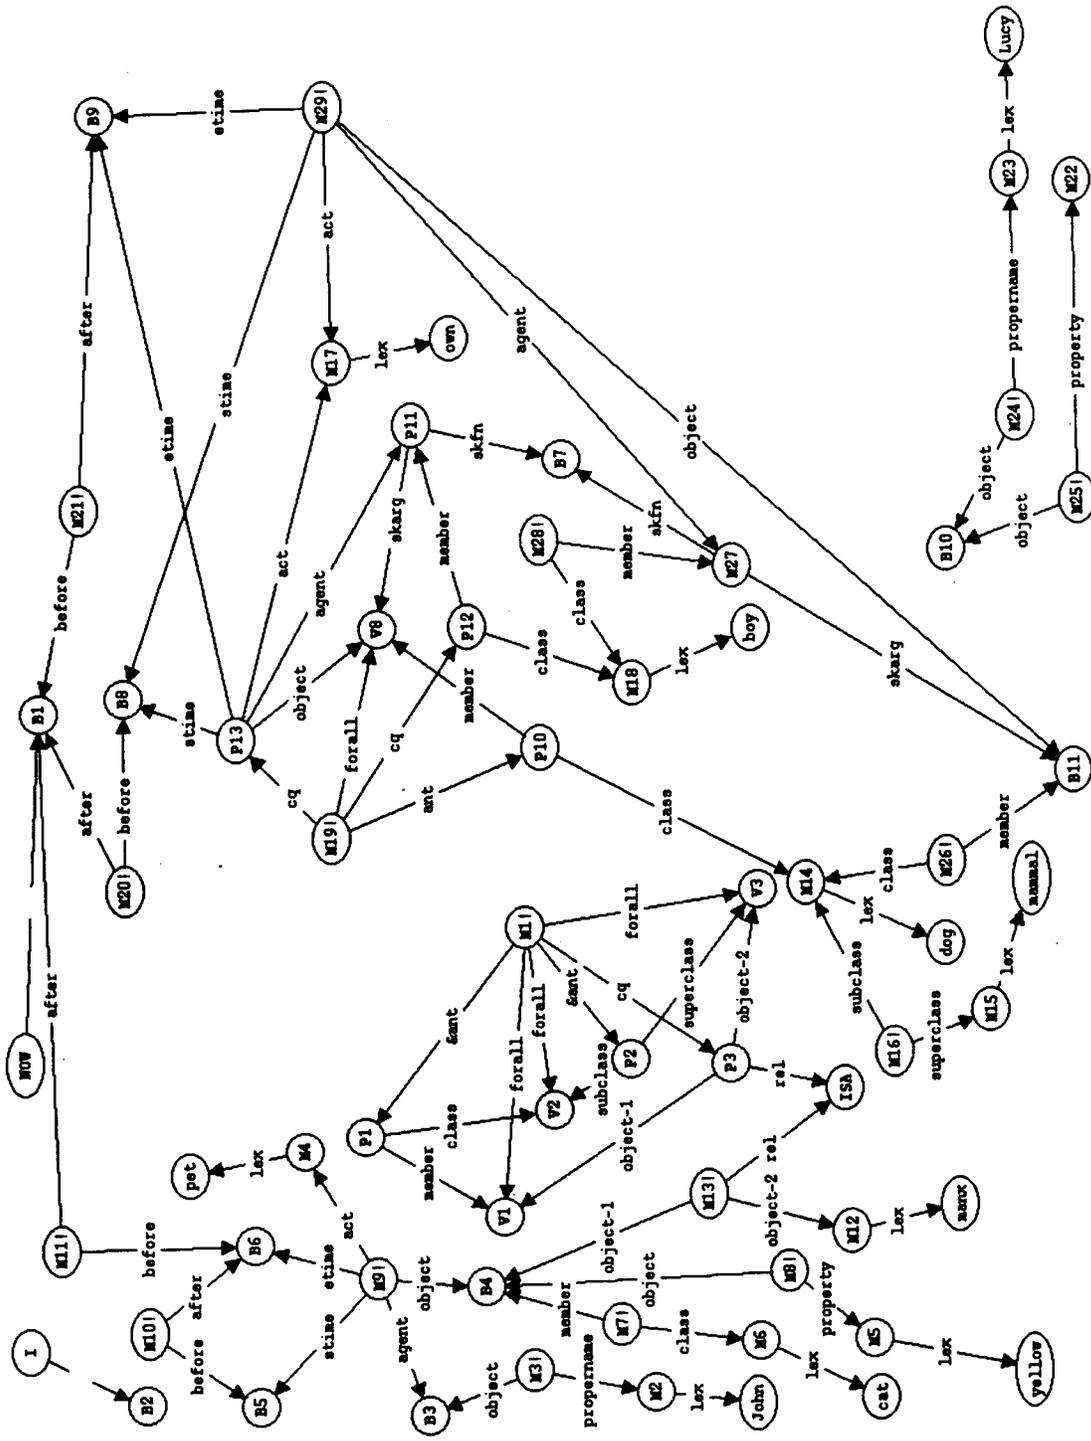
FIG. 4.17. [[M19!]] = Every dog is owned by a boy.

FIG. 4.18. [[B5]] = young Lucy. [[B11]] = a dog.

FIG. 4.19. [[M27]] = the boy.

99

FIG. 4.20. [[M29!]] = The boy owns the dog.

100

FIG. 4.21. [[M31!]] = The dog is a mammal.

FIG. 4.22. [[M32!]] = The dog is yellow.

102

FIG. 4.23. [[M33!]] = Young Lucy petted a yellow dog.

FIG. 4.24. [[M37!]] = The boy saw her.

SNePS networks. Cassie (and SNePS, itself) is still under development, and so can presently read texts only in certain fragments of English. Other uses of Cassie to model readers of narrative texts are discussed in chapters 7 and 8 of this volume.

ACKNOWLEDGMENTS

References for
"An Introduction to a Computational Reader of Narratives"

Stuart C. Shapiro and William J. Rapaport

Department of Computer Science and Engineering

and Center for Cognitive Science

State University of New York at Buffalo, Buffalo, NY 14260-2000

{shapiro | rapaport}@cse.buffalo.edu

http://www.cse.buffalo.edu/sneps

Allaire, Edwin B. (1963), "Bare Particulars," *Philosophical Studies* 14; reprinted in M. J. Loux (ed.), *Universals and Particulars* (Garden City, NY: Anchor Books, 1970): 235–244.

Allaire, Edwin B. (1965), "Another Look at Bare Particulars," *Philosophical Studies* 16; reprinted in M. J. Loux (ed.), *Universals and Particulars* (Garden City, NY: Anchor Books, 1970): 250–257.

Almeida, Michael J. (1987), "Reasoning about the Temporal Structure of Narratives," *Technical Report 87–10* (Buffalo: SUNY Buffalo Department of Computer Science).

Brachman, Ronald J., & Levesque, Hector J. (eds.) (1985), *Readings in Knowledge Representation* (San Mateo, CA: Morgan Kaufmann).

Castañeda, Hector–Neri (1972), "Thinking and the Structure of the World," *Philosophia* 4 (1974) 3–40; reprinted in 1975 in *Critica* 6 (1972) 43–86.

Crimmins, Mark (1989), "Having Ideas and Having the Concept," *Mind and Language* 4: 280–294.

Crimmins, Mark, & Perry, John (1989), "The Prince and the Phone Booth: Reporting Puzzling Beliefs," *Journal of Philosophy* 86: 685–711.

Davidson, Donald (1967), "The Logical Form of Action Sentences," in N. Rescher (ed.), *The Logic of Decision and Action* (Pittsburgh: University of Pittsburgh Press).

Fillmore, Charles (1968), "The Case for Case," in E. Bach & R. T. Harms (eds.), *Universals in Linguistic Theory* (Chicago: Holt, Rinehart & Winston).

Frege, Gottlob (1892), "On Sense and Reference," M. Black (trans.), in P. Geach & M. Black (eds.), *Translations from the Philosophical Writings of Gottlob Frege* (Oxford: Basil Blackwell, 1970): 56–78.

Kumar, Deepak (ed.) (1990), *Current Trends in SNePS—Semantic Network Processing System*, Lecture Notes in Artificial Intelligence, No. 437 (Berlin: Springer-Verlag).

Landman, Fred (1986), "Pegs and Alecs," in J. Y. Halpern (ed.), *Theoretical Aspects of Reasoning About Knowledge* (San Mateo, CA: Morgan Kaufmann): 45–61.

Lehmann, Fritz (ed.) (1992), *Semantic Networks in Artificial Intelligence* (Oxford: Pergamon Press).

Li, Naicong (1986), "Pronoun Resolution in SNePS," *SNeRG Technical Note No. 18* (Buffalo: SUNY Buffalo Department of Computer Science, SNePS Research Group).

1

Maida, Anthony S., & Shapiro, Stuart C. (1982), "Intensional Concepts in Propositional Semantic Networks," *Cognitive Science* 6: 291–330; reprinted in Brachman & Levesque 1985: 169–189.

Martins, João, & Shapiro, Stuart C. (1988), "A Model for Belief Revision," *Artificial Intelligence* 35: 25–79.

Meinong, Alexius (1904), "Über Gegenstandstheorie," in R. Haller (ed.), *Alexius Meinong Gesamtausgabe*, Vol. II (Graz, Austria: Akademische Druck– u. Verlagsanstalt, 1971): 481–535; English translation ("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.

Peters, Sandra L., & Shapiro, Stuart C. (1987a), "A Representation for Natural Category Systems," *Proceedings of the 9th Annual Conference of the Cognitive Science Society (Seattle)* (Hillsdale, NJ: Lawrence Erlbaum Associates): 379–390.

Peters, Sandra L., & Shapiro, Stuart C. (1987b), "A Representation for Natural Category Systems," *Proceedings of the 10th International Joint Conference on Artificial Intelligence (IJCAI-87, Milan)* (San Mateo, CA: Morgan Kaufmann): 140–146.

Peters, Sandra L.; Shapiro, Stuart C.; & Rapaport, William J. (1988), "Flexible Natural Language Processing and Roschian Category Theory," *Proceedings of the 10th Annual Conference of the Cognitive Science Society (Montreal)* (Hillsdale, NJ: Lawrence Erlbaum Associates, 1988): 125–131.

Rapaport, William J. (1978), "Meinongian Theories and a Russellian Paradox," *Noûs* 12: 153–180; errata, *Noûs* 13 (1979) 125.

Rapaport, William J. (1985), "Meinongian Semantics for Propositional Semantic Networks," *Proceedings of the 23rd Annual Meeting of the Association for Computational Linguistics (University of Chicago)* (Morristown, NJ: Association for Computational Linguistics): 43–48.

Rapaport, William J. (1986a), "Searle's Experiments with Thought," *Philosophy of Science* 53: 271–279.

Rapaport, William J. (1986b), "Logical Foundations for Belief Representation," *Cognitive Science* 10: 371–422.

Rapaport, William J. (1988a), "Syntactic Semantics: Foundations of Computational Natural-Language Understanding," in J. Fetzer (ed.), *Aspects of Artificial Intelligence* (Dordrecht, Holland: Kluwer Academic Publishers): 81–131.

Rapaport, William J. (1988b), "To Think or Not to Think," *Noûs* 22: 585–609.

Rapaport, William J. (1991), "Predication, Fiction, and Artificial Intelligence," *Topoi* 10: 79–111.

Rapaport, William J. (1992), "Belief Representation Systems," in S. C. Shapiro (ed.), *Encyclopedia of Artificial Intelligence, Second Edition* (New York: John Wiley & Sons): 98–110.

Rapaport, William J. (forthcoming), "Meinongian Semantics and Artificial Intelligence," in P. Simons (ed.), *Essays on Meinong* (Munich: Philosophia Verlag).

Rapaport, William J.; Shapiro, Stuart C.; & Wiebe, Janyce M. (1986), "Quasi-Indicators, Knowledge Reports, and Discourse," *Technical Report 86-15* (Buffalo: SUNY Buffalo Department of Computer Science).

Rosch, Eleanor (1978), "Principles of Categorization," in E. Rosch & B. B. Lloyd (eds.), *Cognition and Categorization* (Hillsdale, NJ: Lawrence Erlbaum): 27–48.

Routley, Richard (1979), *Exploring Meinong's Jungle and Beyond* (Canberra: Australian National University, Research School of Social Sciences, Department of Philosophy).

Shapiro, Stuart C. (1971), "A Net Structure for Semantic Information Storage, Deduction and

Retrieval," *Proceedings of the Second International Joint Conference on Artificial Intelligence (IJCAI-71; Imperial College, London)* (San Mateo, CA: Morgan Kaufmann), Vol. 2, pp. 512–523.

Shapiro, Stuart 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: Association for Computing Machinery): 219–225.

Shapiro, Stuart C. (1979), "The SNePS Semantic Network Processing System," in N. V. Findler (ed.), *Associative Networks: The Representation and Use of Knowledge by Computers* (New York: Academic Press): 179–203.

Shapiro, Stuart C. (1982), "Generalized Augmented Transition Network Grammars for Generation from Semantic Networks," *American Journal of Computational Linguistics* 8: 12–25.

Shapiro, Stuart C. (1991), "Cables, Paths and 'Subconscious' Reasoning in Propositional Semantic Networks," in J. Sowa (ed.), *Principles of Semantic Networks: Explorations in the Representation of Knowledge* (San Mateo, CA: Morgan Kaufmann): 137–156.

Shapiro, Stuart C. (1993), "Belief Spaces as Sets of Propositions," *Journal of Experimental and Theoretical Artificial Intelligence* 5: 225–235.

Shapiro, Stuart C., & Rapaport, William J. (1987), "SNePS Considered as a Fully Intensional Propositional Semantic Network," in N. Cercone & G. McCalla (eds.), *The Knowledge Frontier: Essays in the Representation of Knowledge* (New York: Springer-Verlag): 262–315; shorter version appeared in *Proceedings of the Fifth National Conference on Artificial Intelligence (AAAI-86, Philadelphia)* (San Mateo, CA: Morgan Kaufmann, 1986): 278–283; revised version reprinted as "SNePS Considered as a Fully Intensional Propositional Semantic Network" in L. Burkholder (ed.), *Philosophy and the Computer* (Boulder, CO: Westview Press, 1992): 75–91.

Shapiro, Stuart C., & Rapaport, William J. (1991), "Models and Minds: Knowledge Representation for Natural-Language Competence," in R. Cummins & J. Pollock (eds.), *Philosophy and AI: Essays at the Interface* (Cambridge, MA: MIT Press): 215–259.

Shapiro, Stuart C., & Rapaport, William J. (1992), "The SNePS Family," *Computers & Mathematics with Applications* 23: 243–275; reprinted in F. Lehmann 1992: 243–275.

Sowa, John (ed.) (1991), *Principles of Semantic Networks: Explorations in the Representation of Knowledge* (San Mateo, CA: Morgan Kaufmann).

Sowa, John (1992), "Semantic Networks," in S. C. Shapiro (ed.), *Encyclopedia of Artificial Intelligence Second Edition* (New York: John Wiley & Sons): 1493–1511.

Srihari, Rohini K. (1981), "Combining Path-Based and Node-Based Inference in SNePS," *Technical Report No. 183* (Buffalo: SUNY Buffalo Department of Computer Science).

Wiebe, Janyce M., & Rapaport, William J. (1986), "Representing *De Re* and *De Dicto* Belief Reports in Discourse and Narrative," *Proceedings of the IEEE* 74: 1405–1413.

Wyatt, Richard (1989), "The Representation of Opaque Contexts," *Technical Report 89-13* (Buffalo: SUNY Buffalo Department of Computer Science).

Wyatt, Richard (1990), "Kinds of Opacity and Their Representations," in D. Kumar (ed.), *Current Trends in SNePS—Semantic Network Processing System*, Lecture Notes in Artificial Intelligence, No. 437 (Berlin: Springer-Verlag): 123–144.

Wyatt, Richard (1993), "Reference and Intensions," *Journal of Experimental and Theoretical Artificial Intelligence* 5: 263–271.