QUASI-INDEXICAL REFERENCE IN PROPOSITIONAL SEMANTIC NETWORKS

William J. Rapaport

Department of Philosophy, SUNY Fredonia, Fredonia, NY 14063
Department of Computer Science, SUNY Buffalo, Buffalo, NY 14260

Stuart C. Shapiro

Department of Computer Science, SUNY Buffalo, Buffalo, NY 14260

ABSTRACT

We discuss how a deductive question-answering system can represent the beliefs or other cognitive states of users, of other (interacting) systems, and of itself. In particular, we examine the representation of first-person beliefs of others (e.g., the _system's_ representation of a _user's_ belief that he himself is rich). Such beliefs have as an essential component "quasi-indexical pronouns" (e.g., 'he himself'), and, hence, require for their analysis a method of representing these pronominal constructions and performing valid inferences with them. The theoretical justification for the approach to be discussed is the representation of _nested_ "_de dicto_" beliefs (e.g., the system's belief that user-1 believes that system-2 believes that user-2 is rich). We discuss a computer implementation of these representations using the Semantic Network Processing System (SNePS) and an ATN parser-generator with a question-answering capability.

1. INTRODUCTION

Consider a deductive knowledge-representation system whose data base contains information about various people (e.g., its users), other (perhaps interacting) systems, or even itself. In order for the system to learn more about these entities--to expand its "knowledge" base--it should contain information about the beliefs (or desires, wants, or other cognitive states) of these entities, and it should be able to reason about them (cf. Moore 1977, Creary 1979, Wilks and Bien 1983, Barnden 1983, and Nilsson 1983: 9). Such a data base constitutes the "knowledge" (more accurately, the _beliefs_) of the system about these entities and about _their_ beliefs.

Among the interrelated issues in knowledge representation that can be raised in such a context are those of multiple reference and the proper treatment of pronouns. For instance, is the person named 'Lucy' whom John believes to be rich the same as the person named 'Lucy' who is believed by the system to be young? How can the system (a) represent the person named 'Lucy' who is an object of its _own_ belief without (b) confusing her with the person named 'Lucy' who is an object of _John's_ belief, yet (c) be able to "merge" its representations of those "two" people if it is later determined that they are the same? A solution to this problem turns out to be a side effect of a solution to a subtler problem in pronominal reference, namely, the proper treatment of pronouns occurring within belief-contexts.

2. QUASI-INDICATORS

Following Castañeda (1967: 85), an _indicator_ is a personal or demonstrative pronoun or adverb used to make a strictly demonstrative reference, and a _quasi-indicator_ is an expression within a 'believes-that' context that represents a use of an indicator by another person. Consider the following statement by person _A_ addressed to person _B_ at time _t_ and place _p_: _A_ says, "I am going to kill you here now." Person _C_, who overheard this, calls the police and says, "_A_ said to _B_ at _p_ at _t_ that he* was going to kill him* there* then*." The starred words are quasi-indicators representing uses by _A_ of the indicators 'I', 'you', 'here', and 'now'. There are two properties (among many others) of quasi-indicators that must be taken into account: (i) They occur only _within_ intentional contexts, and (ii) they cannot be replaced _salva veritate_ by any co-referential expressions.

The general question is: "How can we attribute indexical references to others?" (Castañeda 1980: 794). The specific cases that we are concerned with are exemplified in the following scenario. Suppose that John has just been appointed editor of _Byte_, but that John does not yet know this. Further, suppose that, because of the well-publicized salary accompanying the office of _Byte's_ editor,

(1) John believes that the editor of _Byte_ is rich.

And suppose finally that, because of severe losses in the stock market,

(2) John believes that he himself is not rich.

Suppose that the system had information about each of the following: John's appointment as editor, John's (lack of) knowledge of this appointment, and John's belief about the wealth of the editor. We would _not_ want the system to infer

(3) John believes that he* is rich

because (2) is consistent with the system's information. The 'he himself' in (2) is a quasi-indicator, for (2) is the sentence that _we_ use to express the belief that _John_ would express as 'I am not rich'. Someone pointing to John, saying,

(4) He [i.e., that man there] believes that he* is not rich

could just as well have said (2). The first 'he' in (4) is not a quasi-indicator: It occurs outside the believes-that context, and it can be replaced by 'John' or by 'the editor of Byte', salva veritate. But the 'he*' in)4) and the 'he himself' in (2) could not be thus replaced by 'the editor of Byte' - given our scenario - even though John is the editor of Byte. And if poor John also suffered from amnesia, it could not be replaced by 'John' either.

3. REPRESENTATIONS

Entities such as the Lucy who is the object of John's belief are intentional (mental), hence intensional. (Cf. Frege 1892; Meinong 1904; Castañeda 1972; Rapaport 1978, 1981.) Moreover, the entities represented in the data base are the objects of the system's beliefs, and, so, are also intentional, hence intensional. We represent sentences by means of propositional semantic networks, using the Semantic Network Processing System (SNePS; Shapiro 1979), which treats nodes as representing intensional concepts (cf. Woods 1975, Brachman 1977, Maida and Shapiro 1982).

We claim that in the absence of prior knowledge of co-referentiality, the entities within belief-contexts should be represented separately from entities outside the context that might be co-referential with them. Suppose the system's beliefs include that a person named 'Lucy' is young and that John believes that a (possibly different) person named 'Lucy' is rich. We represent this with the network of Fig. 1.

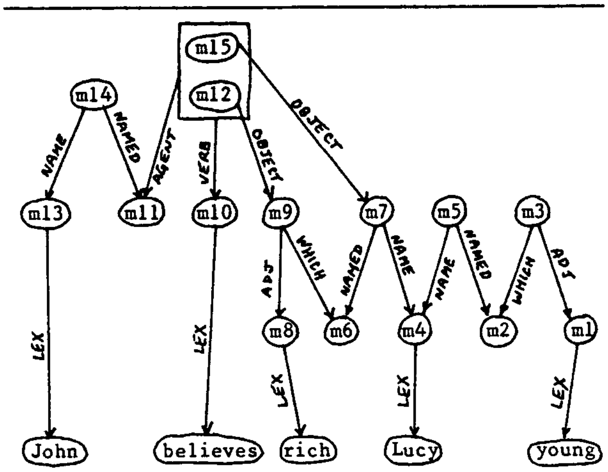

Fig. 1. Lucy is young (m3) and John believes that someone named 'Lucy' is rich (m12).

The section of network dominated by nodes m7 and m9 is the system's de dicto representation of John's belief. That is, m9 is the system's representation of a belief that John might express by 'Lucy is rich', and it is represented as one of John's beliefs. Such nodes are considered as being in the system's representation of John's "belief space". Non-dominated nodes, such as m14, m12, m15, m5, and m3, are the system's representation of its own belief space (i.e., they can be thought of as the object of an implicit 'I believe that' case-frame; cf. Castañeda 1975: 121-22, Kant 1787: B131).

If it is later determined that the "two" Lucies are the same, then a node of co-referentiality would be added (m16, in Fig. 2).

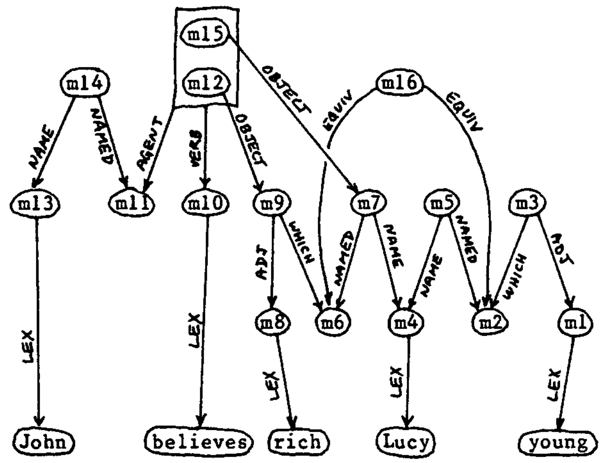

Fig. 2. Lucy is young (m3), John believes that someone named 'Lucy' is rich (m15), and John's Lucy is the system's Lucy (m16).

Now consider the case where the system has no information about the "content" of John's belief, but does have information that John's belief is about the system's Lucy. Thus, whereas John might express his belief as, 'Linus's sister is rich', the system would express it as, '(Lucy system) is believed by John to be rich' (where '(Lucy system)' is the system's Lucy). This is a de re representation of John's belief, and would be represented by node m12 of Figure 3.

The strategy of separating entities in different belief spaces is needed in order to satisfy the two main properties of quasi-indicators.

Consider the possible representations of sentence (3) in Figure 4 (adapted from Maida and Shapiro 1983: 316). This suffers from three major problems. First, it is ambiguous: It could be the representation of (3) or of

(5) John believes that John is rich.

But, as we have seen, (3) and (5) express quite different propositions; thus, they should be separate items in the data base.

Second, Figure 4 cannot represent (5). For then we would have no easy or uniform way to represent (3) in the case where John does not know that he is named 'John': Figure 4 says that the person (m3) who is named 'John' and who believes m6, believes that that person is rich; and this would be false in the amnesia case.

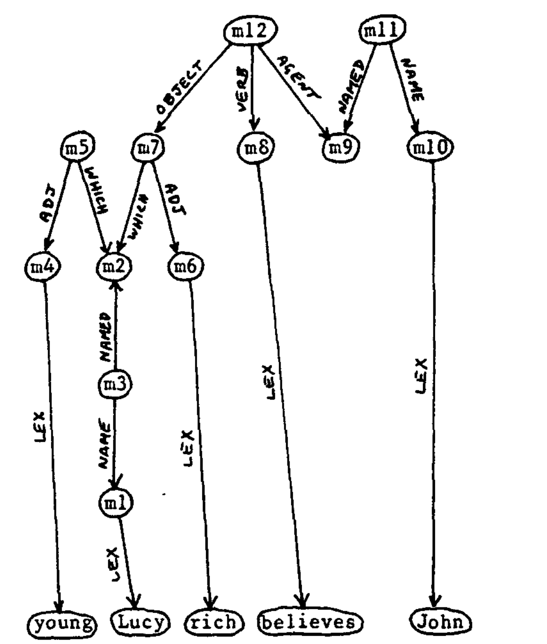

Fig. 3. The system's young Lucy is believed by John to be rich.

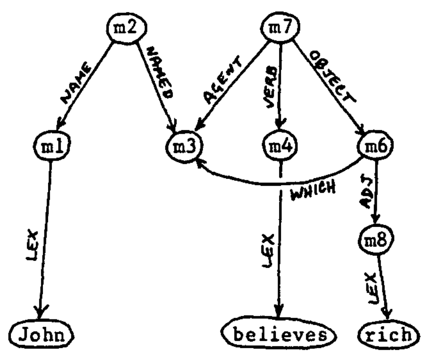

Fig. 4. A representation for 'John believes that he* is rich'

Third, Figure 4 cannot represent (3) either, for it does not adequately represent the quasi-indexical nature of the 'he' in (3): Node m3 represents both 'John' and 'he', hence is both inside and outside the intentional context, contrary to both of the properties of quasi-indicators.

Finally, because of these representational inadequacies, the system would invalidly "infer" (6iii) from (6i)-(6ii):

(6) (i) John believes that he is rich.
(ii) he = John
(iii) John believes that John is rich.

simply because premise (6i) would be represented by the same network as conclusion (6iii).

Rather, the general pattern for representing such sentences is illustrated in Figure 5. The 'he*' in the English sentence is represented by node m2; its quasi-indexical nature is represented by means of node m10.

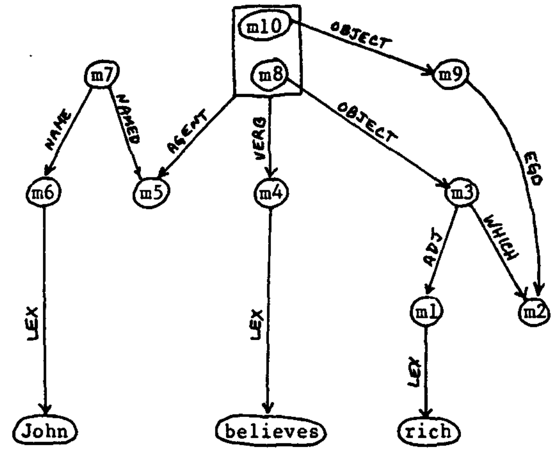

Fig. 5. John believes that he* is rich (m2 is the system's representation of John's "self-concept", expressed by John as 'I' and by the system as 'he*')

That nodes m2 and m5 must be distinct follows from our separation principle. But, since m2 is the system's representation of John's representation of himself, it must be within the system's representation of John's belief space; this is accomplished via nodes m10 and m9, representing John's belief that m2 is his "self-representation". Node m9, with its EGO arc to m2, represents, roughly, the proposition 'm2 is me'.

Our representation of quasi-indexical *de se* sentences is thus a special case of the general schema for *de dicto* representations of belief sentences. When a *de se* sentence is interpreted *de re*, it does not contain quasi-indicators, and can be handled by the general schema for *de re* representations. Thus,

(7) John is believed by himself to be rich

would be represented by the network of Figure 4.

4. INFERENCES

Using an ATN parser-generator with a question-answering capability (based on Shapiro 1982), we are implementing a system that parses English sentences representing beliefs *de re* or *de dicto* into our semantic-network representations, and that generates appropriate sentences from such networks.

It also "recognizes" the invalidity of arguments such as (5) since the premise and conclusion (when interpreted *de dicto*) are no longer represented by the same network. When given an appropriate inference rule, however, the system will treat as valid such inferences as the following:

(8) (i) John believes that the editor of Byte is rich.
(ii) John believes that he* is the editor of Byte.
Therefore, (iii) John believes that he* is rich

In this case, an appropriate inference rule would be:

(9) (∀x,y,z,F)[x believes F(y) & x believes z=y -> x believes F(z)]

In SNePS, inference rules are treated as propositions represented by nodes in the network. Thus, the network for (9) would be built by the SNePS User Language command given in Figure 6 (cf. Shapiro 1979).

```
(build
  avb ($x $y $z $F)
  &ant (build agent *x
              verb (build lex believe)
              object (build which *y
                             adj (build lex *F)))
  &ant (build agent *x
              verb (find lex believe)
              object (build equiv *z equiv *y))
  cq (build agent *x
            verb (find lex believe)
            object (build which *z
                          adj (find lex *F))))
```

Fig. 6. SNePSUL command for building rule (9), for argument (8).

5. ITERATED BELIEF CONTEXTS

Our system can also handle sentences involving iterated belief contexts. Consider

(10) John believes that Mary believes that Lucy is rich.

The interpretation of this that we are most interested in representing treats (10) as the system's de dicto representation of John's de dicto representation of Mary's belief that Lucy is rich. On this interpretation, we need to represent the system's John--(John system)--the system's representation of John's Mary--(Mary John system)--and the system's representation of John's representation of Mary's Lucy--(Lucy Mary John system). This is done by the network of Figure 7.

Such a network is built recursively as follows: The parser maintains a stack of "believers". Each time a belief-sentence is parsed, it is made the object of a belief of the previous believer in the stack. Structure-sharing is used wherever possible. Thus,

(11) John believes that Mary believes that Lucy is sweet

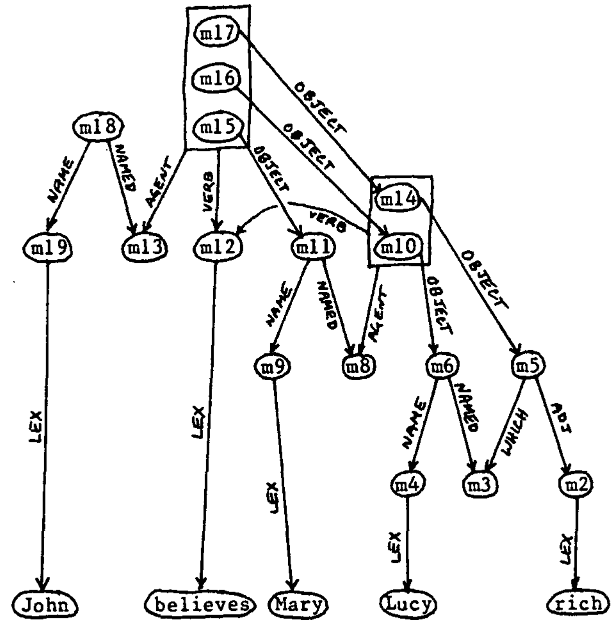

Fig. 7. John believes that Mary believes that Lucy is rich.

would modify the network of Figure 7 by adding new beliefs to (John system)'s belief space and to (Mary John system)'s belief space, but would use the same nodes to represent John, Mary, and Lucy.

6. NEW INFORMATION

The system is also capable of handling sequences of new information. For instance, suppose that the system is given the following information at three successive times:

t1: (12) The system's Lucy believes that Lucy's Lucy is sweet.
t2: (13) The system's Lucy is sweet.
t3: (14) The system's Lucy = Lucy's Lucy.

Then it will build the networks of Figures 8-10, successively. At t1 (Fig. 8), node m3 represents the system's Lucy and m7 represents Lucy's Lucy. At t2 (Fig. 9), m13 is built, representing the system's belief that the system's Lucy (who is not yet believed to be--and, indeed, might not be--Lucy's Lucy) is sweet.[1] At t3 (Fig. 11), m14 is built, representing the system's new belief that there is really only one Lucy. This is a merging of the two "Lucy"-nodes. From now on, all properties of "either" Lucy will be inherited by the "other", by means of an inference rule for the EQUIV case-frame (roughly, the indiscernibility of identicals).

[1]We are assuming that the system's concept of sweetness (node m8) is also the system's concept of (Lucy system)'s concept of sweetness. This assumption seems warranted, since all nodes are in the system's belief space. If the system had reason to believe that its concept of sweetness differed from Lucy's, this could--and would have to--be represented.

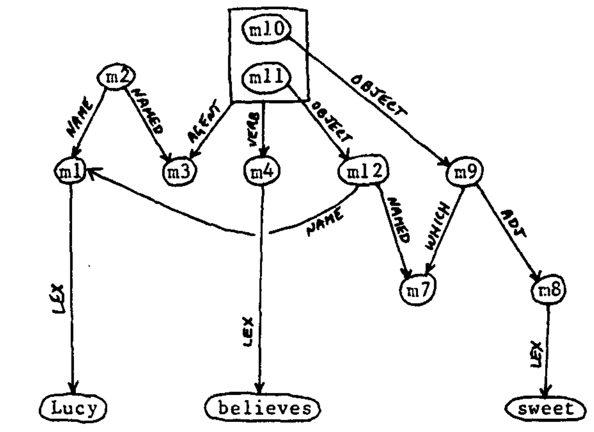

Fig. 8. Lucy believes that Lucy is sweet.

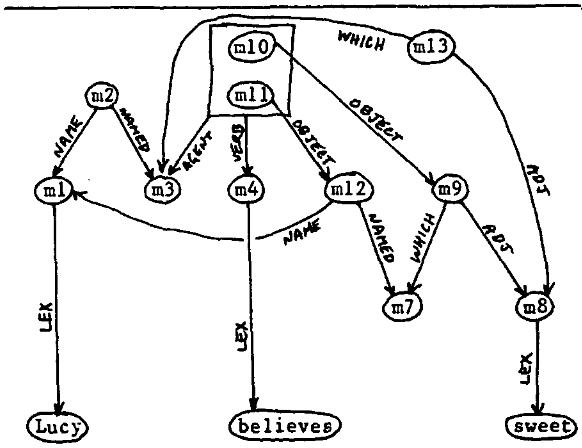

Fig. 9. Lucy believes that Lucy is sweet, and Lucy (the believer) is sweet.

7. FUTURE WORK

There are several directions for future modifications. First, the node-merging mechanism of the EQUIV case-frame with its associated rule needs to be generalized: Its current interpretation is co-referentiality; but if the sequence (12)-(14) were embedded in someone else's belief-space, then co-referentiality might be incorrect. What is needed is a notion of "co-referentiality-within-a-belief-space". The relation of "consociation" (Castañeda 1972) seems to be more appropriate.

Second, the system needs to be much more flexible. Currently, it treats all sentences of the form

(15) x believes that F(y)

as canonically *de dicto* and all sentences of the form

(16) y is believed by x to be F

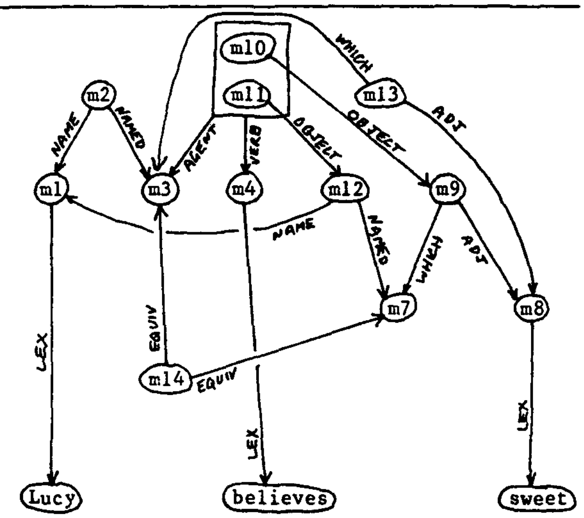

Fig. 10. Lucy believes that Lucy is sweet, Lucy is sweet, and the system's Lucy is Lucy's Lucy.

as canonically *de re*. In ordinary conversation, however, both sentences can be understood in either way, depending on context, including prior beliefs as well as idiosyncracies of particular predicates. For instance, given (1), above, and the fact that John is the editor of *Byte*, most people would infer (3). But given

(17) John believes that all identical twins are conceited.

(18) Unknown to John, he is an identical twin

most people would *not* infer

(19) John believes that he* is conceited.

Thus, we want to allow the system to make the most "reasonable" interpretations (*de re* vs. *de dicto*) of users' belief-reports, based on prior beliefs and on subject matter, and to modify its initial representation as more information is received.

SUMMARY

A deductive knowledge-representation system that is to be able to reason about the beliefs of cognitive agents must have a scheme for representing beliefs. This scheme must be able to distinguish among the "belief spaces" of different agents, as well as be able to handle "nested belief spaces", i.e., second-order beliefs such as the beliefs of one agent about the beliefs of another. We have shown how a scheme for representing beliefs as either *de re* or *de dicto* can distinguish the items in different belief spaces (including nested belief spaces), yet "merge" the items on the basis of new information. This general scheme also enables the system to adequately represent sentences containing quasi-indicators, while not allowing invalid inferences to be drawn from them.

REFERENCES

J. A. Barnden, "Intensions as Such: An Outline," IJCAI-83 (1983)280-286.

R. J. Brachman, "What's in a Concept: Structural Foundations for Semantic Networks," International Journal for Man-Machine Studies 9(1977)127-52.

Hector-Neri Castañeda, "Indicators and Quasi-Indicators," American Philosophical Quarterly 4(1967)85-100.

___, "Thinking and the Structure of the World" (1972), Philosophia 4(1974)3-40.

___, "Identity and Sameness," Philosophia 5(1975)121-50.

___, "Reference, Reality and Perceptual Fields," Proceedings and Addresses of the American Philosophical Association 53(1980)763-823.

L. G. Creary, "Propositional Attitudes: Fregean Representation and Simulative Reasoning," IJCAI-79 (1979)176-81.

Gottlob Frege, "On Sense and Reference" (1892), in Translations from the Philosophical Writings of Gottlob Frege, ed. by P. Geach and M. Black (Oxford: Basil Blackwell, 1970): 56-78.

Immanuel Kant, Critique of Pure Reason, 2nd ed. (1787), trans. N. Kemp Smith (New York: St. Martin's Press, 1929).

Anthony S. Maida and Stuart C. Shapiro, "Intensional Concepts in Propositional Semantic Networks," Cognitive Science 6(1982)291-330.

Alexius Meinong, "Ueber Gegenstandstheorie" (1904), in Alexius Meinong Gesamtausgabe, Vol. II, ed. R. Haller (Graz, Austria: Akademische Druck- u. Verlagsanstalt, 1971): 481-535. English translation in R. Chisholm (ed.), Realism and the Background of Phenomenology (New York: Free Press, 1960): 76-117.

R. C. Moore, "Reasoning about Knowledge and Action," IJCAI-77 (1977)223-27.

Nils J. Nilsson, "Artificial Intelligence Prepares for 2001," AI Magazine 4.4(Winter 1983)7-14.

William J. Rapaport, "Meinongian Theories and a Russellian Paradox," Noûs 12(1978)153-80; errata, 13(1979)125.

___, "How to Make the World Fit Our Language: An Essay in Meinongian Semantics," Grazer Philosophische Studien 14(1981)1-21.

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

___, "Generalized Augmented Transition Network Grammars For Generation From Semantic Networks," American Journal of Computational Linguistics 8(1982)12-25.

Yorick Wilks and Janusz Bien, "Beliefs, Points of View, and Multiple Environments," Cognitive Science 7(1983)95-119.

William A. Woods, "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): 35-79.