QUASI-INDICATORS, KNOWLEDGE REPORTS, AND DISCOURSE

William J. Rapaport, Stuart C. Shapiro,
Janyce M. Wiebe

86-15 June 1986

Department of Computer Science
State University of New York at Buffalo
226 Bell Hall
Buffalo, New York 14260

This work was supported in part by the National Science Foundation under grant no. IST-8504713, and by Research Development Fund Award no. 150-9216-F from the Research Foundation of the State University of New York.

QUASI-INDICATORS, KNOWLEDGE REPORTS, AND DISCOURSE

William J. Rapaport, Stuart C. Shapiro, and Janyce M. Wiebe

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

ABSTRACT

We present a computational analysis of *de re*, *de dicto*, and *de se* belief and knowledge reports. Our analysis solves a problem first observed by Castañeda, namely, that the simple rule '(A knows that P) implies P' apparently does not hold if P contains a quasi-indicator. We present a single rule, in the context of an AI representation and reasoning system, that holds for all propositions P, including quasi-indexical ones. In so doing, we demonstrate the importance of representing proper names explicitly, and we provide support for the necessity of considering sentences in the context of extended text (e.g., discourse or narrative) in order to fully capture certain features of their semantics.

1. INTRODUCTION.

How are knowledge and belief related? The standard philosophical analysis, dating back at least to Plato (*Theaetetus* 201), is that knowledge is justified true belief (but cf. Gettier 1963). In this paper, we describe some issues that are literally in the field of knowledge representation—issues in the representation of knowledge reports, where knowledge is treated as true belief.[1] In particular, we present a computational analysis of *de re*, *de dicto*, and *de se* belief and knowledge reports. Our analysis solves a problem first observed by Castañeda, namely, that the simple rule '(A knows that P) implies P' apparently does not hold if P contains a quasi-indicator.

We present a single rule, in the context of an AI representation and reasoning system, that holds for all propositions P, including quasi-indexical ones. In so doing, we demonstrate the importance of representing proper names explicitly, and we provide support for the necessity of considering sentences in the context of extended text (e.g., discourse or narrative) in order to fully capture certain features of their semantics.

2. DE RE, DE DICTO, AND DE SE BELIEFS.

At the very least, knowledge *implies* true belief and, thus, is a *kind* of belief. Now, among the kinds of belief reports, there are *de re*, *de dicto*, and *de se* belief reports. A *de re* belief report (made by a speaker S to a hearer H), which we shall canonically express as

(1) A believes of N that F,

represents the claim (by S) that agent A believes that someone whom S (and possibly H) believes to be named (or described by) 'N' has property F. Such a report (at least in isolation) is referentially

[1] Consideration of cognitive agents' *justifications* for their beliefs have not recently been of central concern to formal computational analyses of knowledge (cf. Rapaport, forthcoming, for a survey), though, once the appropriate logical foundations for knowledge- and belief-representation are determined, the issue of justification ought once again to become a major area of research.

transparent but propositionally opaque; i.e., '*N*' can be replaced by any co-referring expression, preserving truth value, but at the expense of losing any information about *A*'s characterization of *N*. I.e., '*N*' is a "speaker's reference" and can be replaced by any expression that *S* believes is co-referential with it. (Cf. Castañeda 1970.) E.g., from the *de re* report,

Columbus believed of Castro's island that it was India,

we cannot infer that Columbus characterized that island as being Castro's.

A *de dicto* belief report (by *S* to *H*), which we shall canonically express as

(2) *A* believes that *N* is F,

represents the claim (by *S*) that *A* believes that someone that *he or she* (i.e., *A*) believes to be named (or described by) '*N*' has F; such a report (again, at least in isolation) is referentially opaque but propositionally transparent. I.e., '*N*' is a "believer's reference", and can*not* be replaced by any expression that *S* believes is co-referential. E.g., from the *de dicto* report,

Columbus believed that Queen Isabella was interested in the New World,

we *can* infer that Columbus characterized her as "Queen Isabella", and can*not* replace 'Queen Isabella' by, say, 'the woman described on page 1048 of the *Columbia Encyclopedia*' (even if they *are* co-referential).

Finally, a *de se* belief report (by *S* to *H*) that we shall canonically express as

(3) *A* believes that s/he* is F,

represents a *de dicto* report (by *S* to *H*) involving the *quasi-indicator* 's/he*'. A quasi-indicator is an expression within an intentional context that represents a use of an indicator by another person; indicators, by contrast, make strictly demonstrative reference. (Cf. Castañeda 1966, 1967; Rapaport and Shapiro 1984). Thus, (3) is the *reporter's* (*S*'s) way of expressing the *first-person belief* that *A* would express (using the indicator 'I') as: 'I am F'.

A representation and reasoning system capable of handling these reports in natural language has been implemented using an ATN parser-generator to interface with the SNePS Semantic Network Processing System (Shapiro 1979b, 1982; cf. Rapaport and Shapiro 1984, Rapaport 1984). Figures 1-3 show the formal SNePS representations of these reports.

The analyses we have given for these three types of belief reports can be presented in an informal, linear, predicate language in the following way. We have argued in earlier papers that the terms of an AI representation language should be interpreted as *intensional* entities (in particular, as Meinongian objects), since they are the objects of the "thoughts" of the AI system (Maida and Shapiro 1982, Rapaport 1985, Shapiro and Rapaport 1985). In the following (and later) informal analogues of our SNePS networks, we let Skolem constants m_i range over such intensional entities, and we indicate all predications as in a standard predicate logic:

(F1) Am_2 & Nm_6 & Believe(m_2, Fm_6)
(F2) Am_2 & Believe(m_2, Nm_6) & Believe(m_2, Fm_6)
(F3) Am_2 & Believe(m_2, EGO(m_5)) & Believe(m_2, Fm_5)

The *de re* (F1)—corresponding to Figure 1—says that m_2 is named '*A*', m_6 is named '*N*', and m_2 believes of m_6 that F. The *de dicto* (F2)—corresponding to Figure 2—says that m_2 is named '*A*', m_2 believes of m_6 that m_6 is named '*N*', and m_2 believes of m_6 that F. Note that *de dicto* reports are analyzed in terms of *two de re* reports that are linked (via the common Skolem constant m_6). Finally, the *de se* (F3)—corresponding to Figure 3—says that m_2 is named '*A*', m_2 believes of m_5 that m_5 is him- or herself (i.e., 'EGO(m_5)' is the proposition that m_2 would express as, roughly, 'm_5 is me'), and m_2 believes of m_5 (thus, of him- or herself) that F. Note that (F3) is a *de dicto* report.

Now, just as there are *de re*, *de dicto*, and *de se* belief reports, so, it would seem, there ought to be *de re*, *de dicto*, and *de se* knowledge reports. In this paper, we shall consider to what extent this is so, how various knowledge reports are logically related to their corresponding or underlying belief reports, and the crucial role that extended texts (such as discourse or narrative) play in the analysis.

3. WHAT IS KNOWN IS TRUE.

Since knowledge is *true* belief, epistemic logics (cf. Hintikka 1962) have as a thesis the principle that Barwise and Perry (1983: 196) call "veridicality":

(VK) (A knows that φ) → φ

(where 'A' names a cognitive agent and φ is a proposition. I.e., φ is the proposition that is the objective of A's mental act of knowing, to use the terminology of Rapaport 1985 and Shapiro and Rapaport 1985.) We might express this rule in our system as shown in Figure 4; node m6 represents the rule whose antecedent is node m5 and whose consequent is node φ. (See Shapiro 1979 and Shapiro and Rapaport 1985 for the syntax and semantics of SNePS rules.) Informally, this would be:

Am_2 & $\forall\varphi$[Know(m_2 , φ) → φ]

Prima facie, however, there are three problems with this. First, is such a rule even needed in the system? At first sight, it does not seem to be necessary, since if the system believes that A knows that φ, then surely the system *already* believes that φ and, hence, does not have to *infer* it. But suppose the system comes to believe that A knows that φ because a highly reliable source told it so. It might, then, come to believe that φ by inferring it, using (VK). So let us assume that (VK) should be in the system.

The next two problems that we must face are these: Is (VK) correct for both *de re* and *de dicto* knowledge reports? And, is (VK) correct for *de se* knowledge reports involving quasi-indicators? Let us agree to the following canonical expressions:

(4) A knows that N is F

will express a *de dicto* knowledge report, implying that A believes (*de dicto*) that N is F; and

(5) A knows of N that F

will express a *de re* knowledge report, implying that A believes (*de re*) of N that F.

4. DE RE AND DE DICTO KNOWLEDGE.

Is (VK) correct for both *de re* and *de dicto* knowledge reports? We can split (VK) into two rules, corresponding to the two kinds of reports:

(VK.dd) (A knows that N is F) → (N is F)
(VK.dr) (A knows of N that F) → (N is F)

To express these rules more precisely, we must realize that the belief reports and other propositions represented in the system should be treated as beliefs of the system. I.e., the system should be treated as a cognitive agent. We have named the cognitive agent implemented by our system, 'Cassie' (for *C*ognitive *A*gent of the *S*NePS *S*ystem—an *I*ntelligent *E*ntity).

Thus, in the *de dicto* case, we can express the thesis that knowledge is true belief as follows:

(KTB.dd) (Cassie believes that A knows that N is F)
→ (Cassie believes that A believes that N is F) & (Cassie believes that N is F)

Note that now it is no longer the case that knowledge implies true belief *simpliciter*; rather, Cassie's belief about a knowledge report (the antecedent of (KTB.dd) implies her belief about a belief report (the first conjunct of the consequent of (KTB.dd) and her belief about the objective of that report (the second conjunct). The consequent of (KTB.dd) trivially implies that Cassie believes that N is F, agreeing with (VK.dd).

In the *de re* case, the knowledge-is-true-belief thesis becomes:

(KTB.dr) (Cassie believes that A knows of N that F)
→ (Cassie believes that A believes of N that F) & (Cassie believes that N is F)

whose consequent again trivially implies that Cassie believes that N is F, agreeing with (VK.dr).

Before proceeding, it is important to get clear on a central point. In the *de dicto* case, (KTB.dd), *N* must be "in" *A*'s "belief space". I.e., '*N*' is *A*'s characterization of the individual that is F. In the *de re* case, (KTB.dr), *N* must be directly in Cassie's belief space, but is not necessarily in *A*'s. I.e., '*N*' is Cassie's characterization of the individual that is F. (Of course, all nodes are trivially in Cassie's belief space. E.g., in the *de dicto* case, '*N*' is really Cassie's characterization of *A*'s characterization of the individual. Cf. Wiebe and Rapaport (forthcoming) for details.)

SNePS representations of these rules are shown in Figures 5-6. In our informal, intensional, predicate notation, these become:

(F5) $\text{Know}(m_2, \text{F}m_6) \rightarrow \text{F}m_6$

where $\text{A}m_2$ and $\text{N}m_6$, and

(F6) $\text{Know}(m_2, \text{F}m_6) \,\&\, \text{Know}(m_2, \text{N}m_6) \rightarrow \text{F}m_6 \,\&\, \text{N}m_6$

where $\text{A}m_2$.

It is important to note that (KTB.dd) (in Fig. 6) is redundant (in the presence of (KTB.dr) in Fig. 5): Since our analysis of *de dicto* reports is essentially a conjunction of two, linked, *de re* reports, two applications of (KTB.dr)—to nodes m10 and m11 of Figure 6—yield both consequents of (KTB.dd). In other words, (F6) is a conjunction of two instances of the general form of (F5):

(F5G) $\forall \text{P} \forall m \forall n [\text{Know}(m, \text{P}n) \rightarrow \text{P}n]$

We repeat, for emphasis, that *m* and *n* here do not range over ***names*** of individuals, but over *concepts* of individuals, who may or may not be named or otherwise described.

5. DE SE KNOWLEDGE.

The veridicality thesis does *not* hold when the objective contains a quasi-indicator (Castañeda 1966, 1967). This can be seen in the general case (we use ' * ' instead of the more awkward 'he* or she*'):

(VK.*) (*A* knows that * is F) → (* is F)

cannot be true, since the occurrence of the quasi-indicator ' * ' in the *consequent* is not within the scope of an intentional verb, and, hence, it has no antecedent: we cannot simply detach the consequent, since it cannot stand by itself, so to speak. It is even easier to see this if we bring Cassie into the picture. In the case of a *de dicto/de se* knowledge report—which, because it is *de dicto*, involves a quasi-indicator—we have:

(KTB.ds) (Cassie believes that *A* knows that * is F)
→ (Cassie believes that *A* believes that * is F) & (Cassie believes that she* is F).

The SNePS representation of part of this rule is shown in Figure 7. Informally, it is (where $\text{A}m_2$):

(F7) $\text{Know}(m_2, \text{F}m_5) \,\&\, \text{Know}(m_2, \text{EGO}(m_5)) \rightarrow \text{F}m_5 \,\&\, \text{EGO}(m_5)$.

The SNePS Inference Package will assert the propositions labelled m8 and m6 (i.e., the consequents of (F7)), thus representing—incorrectly—that Cassie believes that ***she**** is F. Note, again, that (KTB.ds) is redundant: two applications of (KTB.dr)—to nodes m9 and m10 of Figure 7—yield both consequents of (KTB.ds).

Clearly, what we would like is not (KTB.ds), but

(KTB.ds.1) (Cassie believes that *A* knows that * is F)
→ (Cassie believes that *A* believes that * is F) & (Cassie believes that *A* is F).

part of which can be represented in SNePS as in Figure 8. Informally (where $\text{A}m_2$):

(F8) $\text{Know}(m_2, \text{F}m_5) \,\&\, \text{Know}(m_2, \text{EGO}(m_5)) \rightarrow \text{F}m_2$.

To emphasize that this is the only troubling case, consider a *de re/de se* knowledge report—which, because it is *de re*, does *not* involve a quasi-indicator (Rapaport 1984, Sect. V)—we have:

(KTB.drds) (Cassie believes that *A* knows of him/herself that F)
→ (Cassie believes that *A* believes of him/herself that F) & (Cassie believes that *A* is F).

In this non-quasi-indexical, *de re/de se* case, we have the same consequent as in the non-quasi-indexical, non-*de se*, *de re* case (KTB.dr): For the antecedent of (KTB.drds) is equivalent (by referential transparency) to: Cassie believes that *A* knows of *A* that F. (See Fig. 9; informally: Know(m_2, Fm_2) → Fm_2, where Am_2.) But the consequent of (KTB.ds)—the quasi-indexical, *de dicto/de se* case—is *not* the same as in the non-quasi-indexical, *de dicto* case (KTB.dd). In the former, Cassie believes that *A* believes that * is F; in the latter, Cassie believes that *A* believes that *A* is F.

The main problem is this: it will not suffice to have a separate rule, namely (KTB.ds.1), for the quasi-indexical case, since the rule for the *de re* case (KTB.dr)—and hence the rule for the *de dicto* case, (KTB.dd)—will still allow the inference that we don't want; i.e., (KTB.ds)—which is what we don't want—is just a special case of (KTB.dd) and, hence, of (KTB.dr).

6. A SOLUTION.

The broader context of our problem is this: In earlier work (Rapaport and Shapiro 1984, Rapaport 1984), we argued that quasi-indexical reference must be capable of being handled by a belief-representation system, and we presented a computationally adequate mechanism for doing this. That mechanism was adequate as long as we only considered belief reports *in isolation*. When we turn to embedded text, where conjunctions—especially *sequences*—of belief reports are considered—as in discourse or narrative—the data become more complex, and a correspondingly more complex theory is needed. In Wiebe and Rapaport (forthcoming), we show that when such sequences are considered, the notions of referential and propositional opacity and transparency interact in ways that blur the distinctions among them. In this paper, we show that our original representation of quasi-indicators must be modified in order to handle knowledge reports, which are, in fact, conjunctions of belief reports.

The solution we now propose is to represent quasi-indexical, *de se/de dicto* belief and knowledge reports as shown in Figure 10. Informally,

(F10) Am_2 & Know(m_2, Fm_2).

Notice that there is no "EGO belief" component, as in (F3). Using this representation, the inference from

Cassie believes that *A* knows that * is F

to

Cassie believes that *A* is F

can be handled by the *same* rule (KTB.dr) as in the other cases (roughly because Am_2 is outside the scope of 'Know'), and—because there is no "EGO belief" component—'Cassie believes that she* is F' is no longer inferrable.

However, there are several potential problems that must be cleared up before this solution can be adopted. First, Figure 10 is a representation for quasi-indicators that was *rejected* in our earlier work ! So, we must re-examine those arguments. Second, the representation in Figure 10 does not appear to be *de dicto* (since it does not consist of *two*, linked, *de re* belief reports); so we must re-examine the nature of *de dicto* belief reports to see whether our claim that quasi-indexical belief is *de dicto* can be maintained. Third, our original representation made use of an EGO arc and a representation of *A*'s "self-concept", whereas our new representation does not. But the notion of an agent's self-concept is of independent importance, so we must explore alternative representations for it. We now turn to an exploration of these issues.

6.1. Is Figure 10 Acceptable?

6.1.1. Is Figure 10 Ambiguous?

In Rapaport and Shapiro 1984 and Rapaport 1984, we rejected the representation of Figure 10 on the grounds that it ambiguously represented both

A believes that * is F

and

A believes that *A* is F,

which are not equivalent. But the latter really should be represented as in Figure 11. Informally,

(F11) Am_2 & Believe(m_2 , Fm_5) & Believe(m_2 , Am_5).

So, the representation of Figure 10 *is* available to represent the former. The Figure-10 representation is ambiguous only if m_2 is interpreted as a name, which we do not do. This issue is taken up in Section 6.1.3.

6.1.2. Is Figure 10 Quasi-Indexical?

We also argued that the Figure-10 network did not adequately represent the quasi-indexical nature of the belief report, on the grounds that node m2—representing *A*'s self-concept—was both inside and outside the intentional context—i.e., in both Cassie's and *A*'s belief spaces. But, of course, *all* nodes are in Cassie's belief space, and what must be represented is *Cassie's* belief, which is that the person believed by *A* to be F is *A*—the believer—himself or herself. Figure 10 does represent this; what it does not—and should not—do is suppose that *A* characterizes him- or herself with the *name* '*A*'.

6.1.3. The Proper Treatment of Proper Names.

The original motivation, however, for the Figure-3 representation was *not* the alleged ambiguity of Figure 10, but the *actual* ambiguity of Maida and Shapiro's representation (1982) shown in Figure 12. Here, it should be noted, the PROPER-NAME–OBJECT case frame is not used. Informally,

(F12) Believe(*A*, F*A*).

Note that here '*A*' is the Skolem constant; it is *not* a proper name.

The proper lesson to be learned from this is the importance of the PROPER-NAME–OBJECT proposition for the representation of cognitive agents. Shapiro used such propositions before the Maida and Shapiro paper (using a NAME-NAMED case frame; Shapiro 1975, 1979, 1982), but felt that nothing major was lost by abbreviating the representation used in Maida and Shapiro 1982 to the extent of not separately showing this proposition. It was the abbreviated version that Rapaport realized was ambiguous between the *de re* and *de dicto* cases, and this led us to the EGO proposition (in Rapaport and Shapiro 1984 and Rapaport 1984). We now see that, although the EGO proposition works when representing nested beliefs, it does not work when representing nested knowledge.

The lesson is: When representing a cognitive agent within a belief system, it is important to represent the agent in a way that is neutral to any properties (including its name) ascribed to it by the believer. In that way, the representation of the cognitive agent may be used in representations of its beliefs about itself without automatically ascribing to it any of the properties ascribed to it by the believer. If the representation is not neutral, and the automatic transfer of the property ascription is not wanted, node splitting must be used (see Maida and Shapiro 1982).

6.2. How to Represent a Self-Concept.

With the EGO arc, we are able to represent Cassie's beliefs about herself. It is essential that we be able to do this. Not only must we be able to represent Cassie's belief, say, 'I am intelligent', but Cassie might have false nested beliefs about herself or fail to believe that she in fact has certain beliefs about herself. E.g., Cassie might explicitly believe that she believes that φ, yet she might not in fact believe that φ (as evidenced by her failure to act in accordance with φ). Or Cassie might in fact believe that φ, yet not believe that she believes it (or, of course, believe that she does not believe it).

Without the EGO arc, how can we represent these? The solution we have chosen is a generalization of a mechanism that our research group uses for representing the temporal indexical 'now':

namely, a node representing 'now' is identified by a (movable) 'now'-pointer. The "temporal" node pointed to by the 'now'-pointer will change as linguistic cues in the discourse or narrative move the 'now' point along (cf. Almeida and Shapiro 1983, 1986; Shapiro and Rapaport 1985; Almeida 1986).

Similarly, within Cassie's belief space, we postulate an 'I'-pointer, which, at the beginning of a dialogue with Cassie, is initialized to point to a node, which will then represent Cassie's self-concept. Unlike the 'now'-pointer, the 'I'-pointer does not need to be updated. On the other hand, just as, when reading a narrative, 'now'-points are stacked when entering sub-narratives (e.g., a flashback), the 'I'-pointer is stacked when entering nested belief spaces. At the top level, the word 'I' is used to express the node pointed to by the 'I'-pointer; when the context is a nested belief space, the word 'I' would change to 'she*' or 'he*' (as appropriate).

6.3. Is Figure 10 De Dicto?

Quasi-indexical *de se* beliefs are *de dicto*. This is, perhaps, arguable. But, like *de dicto* and unlike *de re* belief reports, they are referentially opaque and propositionally transparent, at least in isolation. Yet Figure 10 does not have the structure of a *de dicto* report; indeed, it appears to have the structure of a (single) *de re* report.

Now, the *de dicto/de se* report

(6) *A* believes that * is F

implies, but is not implied by, the *de re/de se* report

(7) *A* believes of him/herself that F.

Figure 10 is the representation of (6); it is also *a* representation of (7), which is consistent with the fact that (6) implies (7). But in various contexts, various representations will be used to represent (7) (e.g., Figs. 10, 13, etc.). So it is not the case *in general* that (7) implies (6).

6.3.1. Castañeda-Style Predication.

Is there, though, a way to represent the quasi-indexical *de se* belief in such a way that it wears its *de-dicto*-ness on its sleeve, so to speak? There is, but it might be otiose. Our (F1)-analysis of *de dicto* belief reports is this:

A believes that *N* is F

is analyzed as (a Skolemized form of):

A believes that something that is named '*N*' is (the same as something that is) F.

Similarly, our (F3)-analysis of

A believes that * is F

is (a Skolemized form of):

A believes that something that is * is (the same as something that is) F

These suggest the patently *de dicto* SNePS networks of Figures 14 and 15. The mode of predication exhibited here is not a simple OBJECT-PROPERTY case frame. Rather, '*N* is F' is analyzed as (a Skolemized form of):

$\exists x$ [x is named '*N*' & x is F].

This is very close to the theory of predication put forth in Castañeda 1972 (where the Skolem constants would now be interpreted as ranging over "guises", and which we previously urged as an analysis of predication in SNePS (Rapaport 1985). It can now be seen to have the additional advantage of exhibiting the *de dicto* nature of quasi-indexical *de se* reports.

Does it run into the same problem that our earlier *de dicto/de se* representation does with respect to knowledge? No; 'Cassie believes that *A* knows that * is F' would simply imply that Cassie believes that someone is F and that that someone is *A*, which is precisely right.

So, is the extra belief about the equivalence of the object that is F and the object named '*N*' needed? If not, then the representation of Figure 10 suffices (at least till more complex data is unearthed). We think that it is *not* needed, at least in order to render the Figure-10 analysis *de dicto*. But to show this, we advocate a new understanding of *de dicto* and *de re* belief reports in the context of discourse.

6.3.2. A New Theory of De Re and De Dicto Belief Reports.

Consider two participants in a dialogue, Cassie and Oscar (the *O*ther *S*NePS *C*ognitive *A*gent *R*epresentation). Suppose that Oscar says to Cassie (perhaps in a vain attempt to impress her),

I am rich,

thus expressing the belief represented in Figure 16A. Cassie's *interpretation* of this is expressed by her as

Oscar believes that he* is rich

and represented (using the Figure-10 representation) as in Figure 16B. Suppose, next, that Oscar says to Cassie (perhaps in a vain attempt to make her jealous),

Lucy is sweet,

thus expressing the belief represented in Figure 17A. Cassie's *interpretation* of this is expressed by her as

Oscar believes that Lucy is sweet

and represented as in Figure 17B.

I.e., representations of *de dicto* belief reports are *Cassie's interpretations* of reports made *by* the believer (i.e., reports *from* the believer *to* Cassie *about* him- or herself), and are such that Cassie's representation is "exactly" like the believer's representation, except for (1) the fact that all nodes are in Cassie's belief space, not the believer's, and (2) the shift from *indicators* (used by the believer) to *quasi*-indicators (used by Cassie), which is represented by the use of an embedding belief-structure in Cassie's belief space in place of the 'I'-pointer in Oscar's belief space. That is, Oscar's 'I'-point becomes Cassie's 'Oscar'-point, so to speak.

Finally, suppose that a third person, Boris, knows that Oscar's Sue is Cassie's Mary (i.e., that the person Oscar believes to be Sue is the person Cassie believes to be Mary) and that Boris tells Cassie that Oscar believes of Mary that she is tall. I.e., Boris believes that Oscar believes that Sue is tall, represented in Figure 18A. Cassie's *interpretation* of this third-person, *de re* report is that Oscar believes of Mary that she is tall, represented in Figure 18B. I.e., *de re* belief reports are *Cassie's interpretation* of a *third* person's interpretation of Oscar's beliefs (i.e., reports *from* a third person *to* Cassie *about* the believer), and are such that Cassie's representation is like Oscar's only with respect to the fragment that is in common. This is the core of what is meant by 'propositional opacity'. (*De re* reports might also be *inferred* by Cassie from other beliefs that she has.)

There is one final issue to consider. Suppose that Cassie is told by Boris that Oscar believes of the person who Cassie and Boris believe is Oscar that he is rich. Should Cassie interpret this as in Figure 19 or Figure 20 (cf. Figs. 10, 13, respectively)? If Cassie interprets Boris's belief report as in Figure 19, then she could infer that Oscar believes that he* is rich, which might be false. So Figure 20 ought to be Cassie's interpretation. If Boris then tells Cassie that Oscar believes that he* is rich, then Figure 20 would be modified as in Figure 21, because Cassie still does not know whether Oscar believes that two people or one person is rich. Finally, if Boris tells Cassie that Oscar only believes himself* to be rich, then Cassie must "merge" two nodes in her representation of Oscar's belief space, as in Figure 22 (cf. Maida and Shapiro 1982).

7. CONCLUSION.

There are several points that we have tried to make in this paper. The first is that the simple rule

(VK) does not always hold; this is the negative point first made by Castañeda some 20 years ago but not hitherto incorporated in computational analyses of knowledge and belief. The second is a positive contribution: a single rule, implementable in SNePS, that can replace (VK)—namely (KTB.dr) (Fig. 5; cf. (F5G)). Third, we demonstrated the importance of representing proper names explicitly. Fourth, we provided support for the necessity of considering sentences in the context of extended text in order to fully capture certain features of their semantics. [2]

REFERENCES

(1) Almeida, Michael J. (forthcoming), "Reasoning about the Temporal Structure of Narratives," SUNY Buffalo Dept. of Computer Science Technical Report.

(2) Almeida, Michael J., and Shapiro Stuart C. (1983), "Reasoning about the Temporal Structure of Narrative Texts," *Proc. 5th Annual Meeting Cognitive Science Soc.* (Rochester, NY).

(3) Almeida, Michael J., and Shapiro, Stuart C. (forthcoming), "The Construction of the Narrative Line."

(4) Barwise, Jon, and Perry, John (1983), *Situations and Attitudes* (Cambridge, Mass.: MIT Press).

(5) Castañeda, Hector-Neri (1966), " 'He': A Study in the Logic of Self-Consciousness," *Ratio* 8: 130-57.

(6) Castañeda, Hector-Neri (1967), "Indicators and Quasi-Indicators," *Amer. Phil. Qtly.* 4: 85-100.

(7) Castañeda, Hector-Neri (1970), "On the Philosophical Foundations of the Theory of Communication: Reference," *Midwest Studies in Phil.* 2(1977)165-86.

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

(9) Gettier, Edmund L. (1963), "Is Justified True Belief Knowledge?", *Analysis* 23: 121-23; reprinted in A. P. Griffiths (ed.), *Knowledge and Belief* (Oxford: Oxford Univ. Press, 1967).

(10) Hintikka, Jaakko (1962), *Knowledge and Belief: An Introduction to the Logic of the Two Notions* (Ithaca: Cornell Univ. Press).

(11) Maida, Anthony S., and Shapiro, Stuart C. (1982), "Intensional Concepts in Propositional Semantic Networks," *Cognitive Science* 6: 291-330; reprinted in R. J. Brachman and H. J. Levesque (eds.), *Readings in Knowledge Representation* (Los Altos, CA: Morgan Kaufmann, 1985): 169-89.

(12) Plato, *Theaetetus*, in E. Hamilton and H. Cairns (eds.), *The Collected Dialogues of Plato* (Princeton: Princeton Univ. Press, 1961): 845-919.

(13) Rapaport, William J. (1984), *Belief Representation and Quasi-Indicators.* SUNY Buffalo Dept. of Computer Science Technical Report 214.

(14) Rapaport, William J. (1985), "Meinongian Semantics for Propositional Semantic Networks," *Proc. 23rd Annual Meeting Assoc. for Computational Linguistics* (Univ. of Chicago): 43-48.

(15) Rapaport, William J. (forthcoming), "Belief Systems," in S. C. Shapiro (ed.), *Encyclopedia of Artificial Intelligence* (New York: Wiley).

(16) Rapaport, William J., and Shapiro, Stuart C. (1984), "Quasi-Indexical Reference in Propositional Semantic Networks," *Proc. 10th Int'l. Conf. Computational Linguistics* (COLING-84) (Stanford Univ.): 65-70.

---

[2] This material is based upon work supported by the National Science Foundation under Grant No. IST-8504713 and by Research Development Fund Award #150-9216-F from the Research Foundation of State University of New York. We are grateful to our colleagues in the SNePS Research Group for discussions on these issues.

(17) Shapiro, Stuart C. (1975), "Generation as Parsing from a Network into a Linear String," *American J. Computational Linguistics*, Microfiche 33: 45-62.

(18) Shapiro, Stuart C. (1979a), "Generalized Augmented Transition Network Grammars for Generation from Semantic Networks," *Proc. 17th Annual Meeting Assoc. for Computational Linguistics* (Univ. of California at San Diego): 25-29.

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

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

(21) Shapiro, Stuart C., and Rapaport, William J. (1985) "SNePS Considered as a Fully Intensional Propositional Semantic Network," SUNY Buffalo Dept. of Computer Science Technical Report 85-15; forthcoming in G. McCalla and N. Cercone (eds.), *Knowledge Representation* (Berlin: Springer-Verlag, forthcoming).

(22) Wiebe, Janyce M., and Rapaport, William J. (forthcoming), "Understanding *De Re* and *De Dicto* Belief Reports in Discourse and Narrative," *Proc. IEEE*, Special Issue on Knowledge Representation.

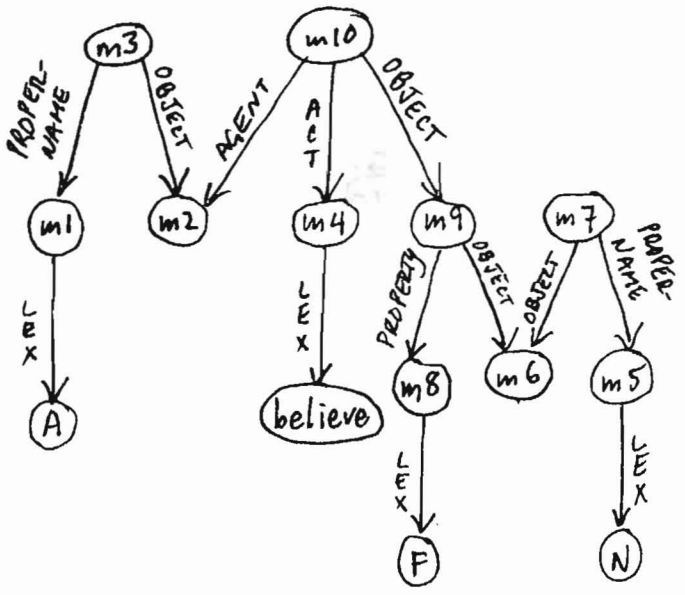

Fig. 1: A believes (de re) of N that F

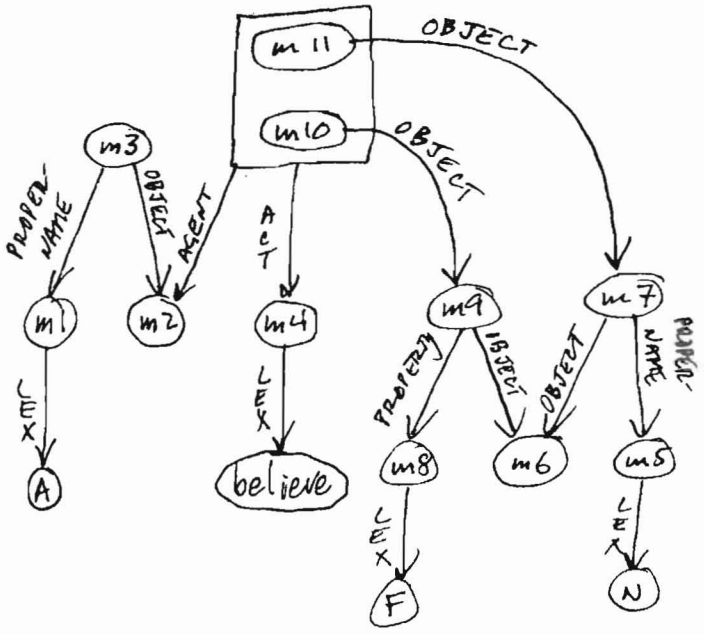

Fig. 2: A believes (de dicto) that N is F

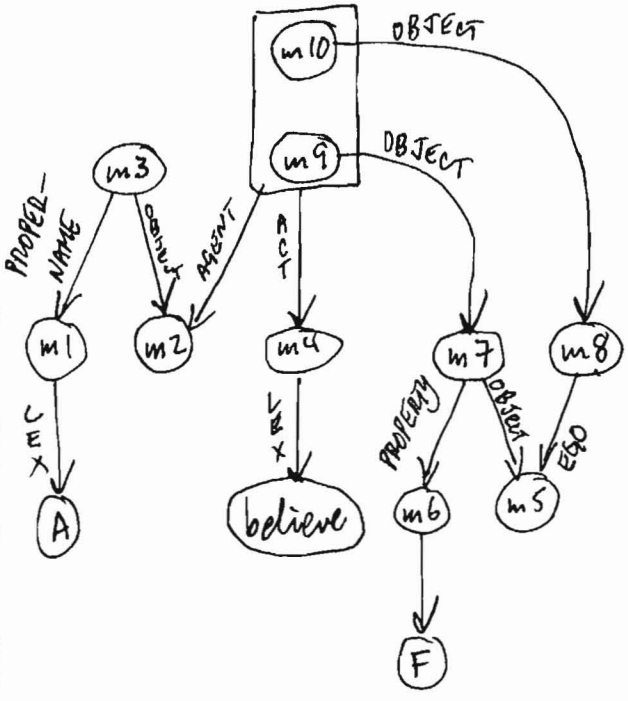

Fig. 3: A believes (de dicto) that s/he* is F
(a de se belief report)

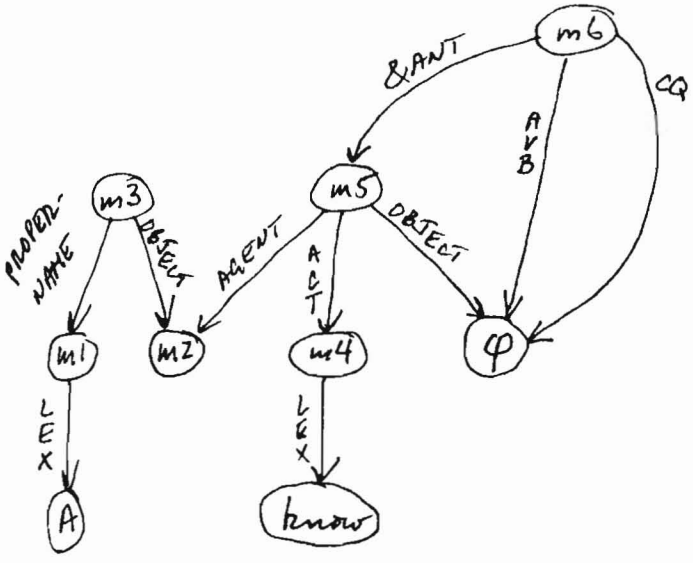

Fig. 4: (A knows that ϕ) $\rightarrow \phi$

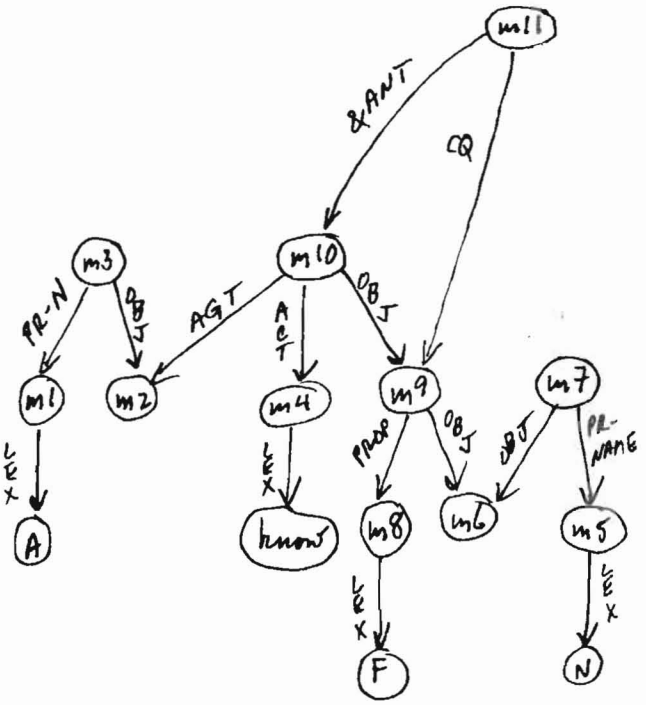

Fig. 5: (A knows of N that F) → N is F

(N.B.: the AVB quantifier arcs are not shown)

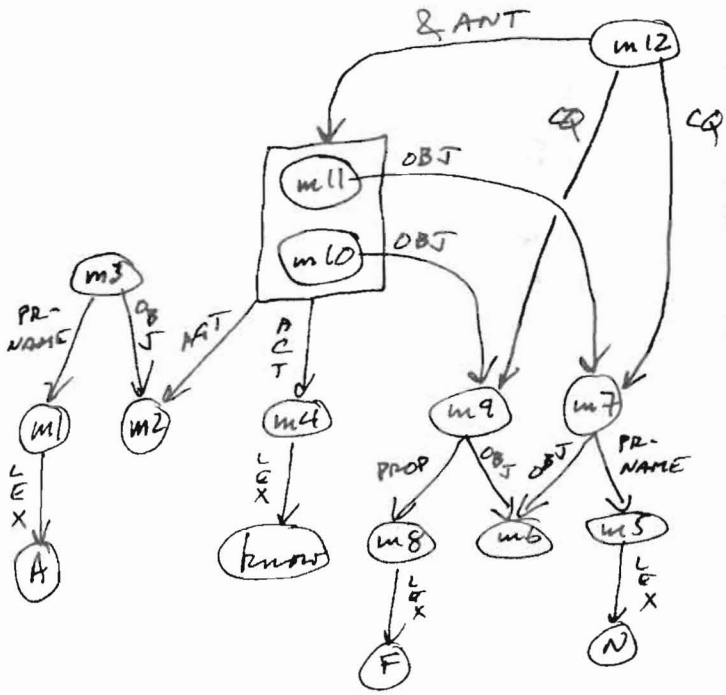

Fig. 6: (A knows that N is F) → N is F

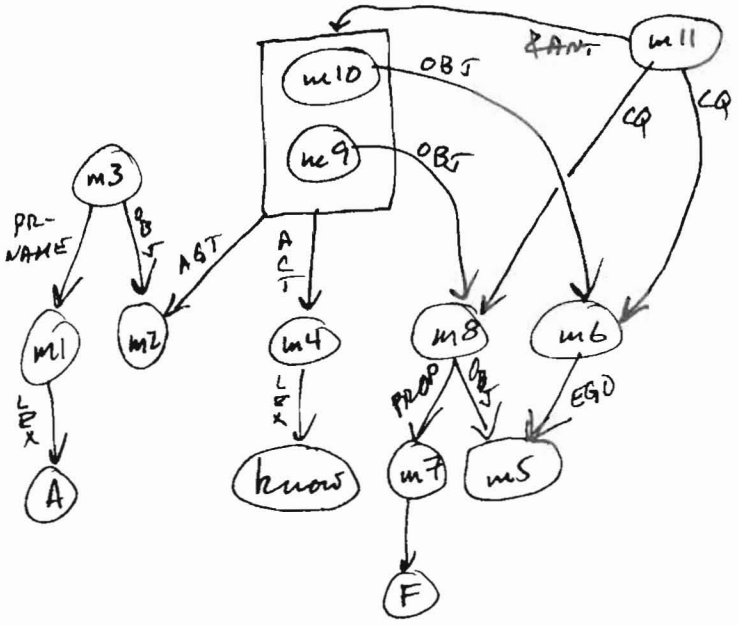

Fig. 7: (A knows that * is F) → (* is F)

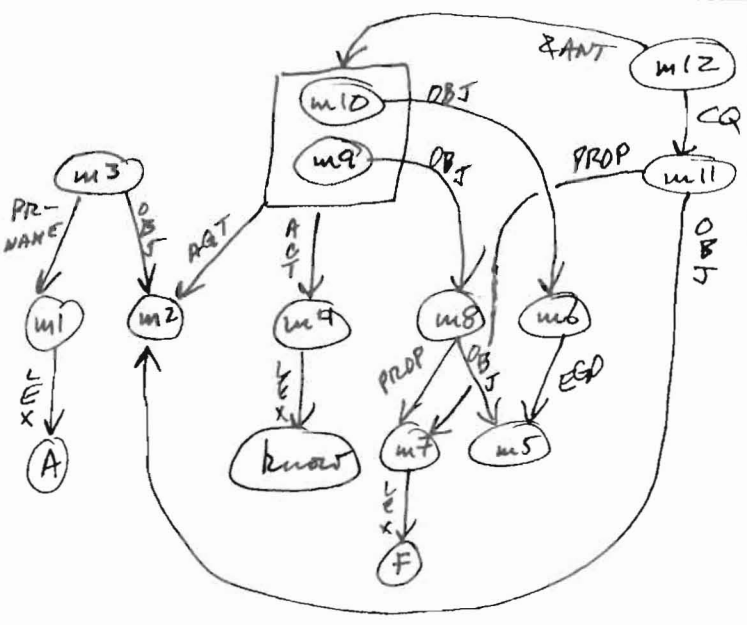

Fig. 8: (A knows that * is F) → (A is F)

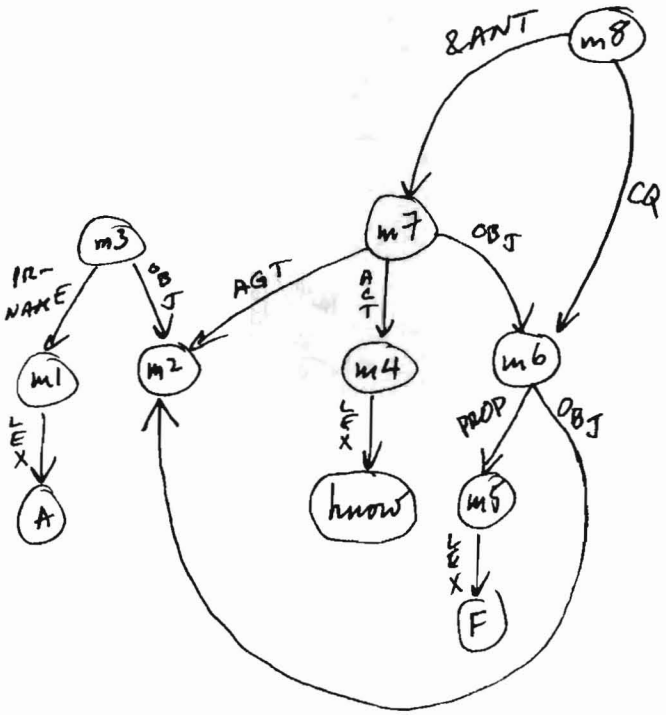

Fig. 9: (A knows of A that F) → (A is F)

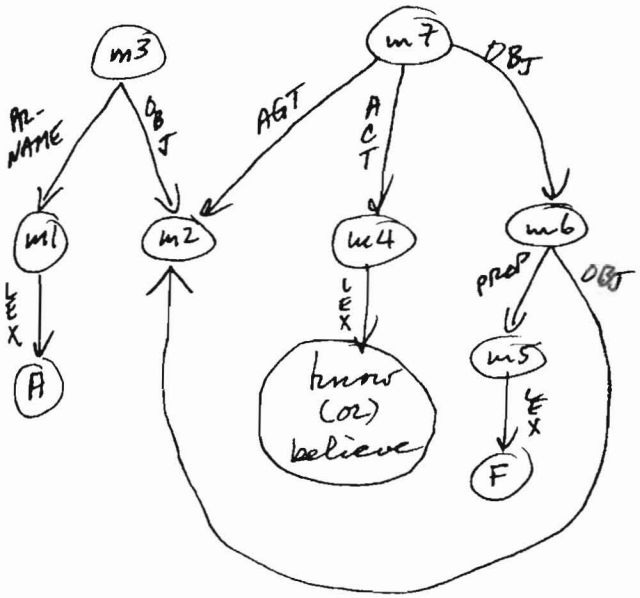

Fig. 10: A believes/knows that * is F

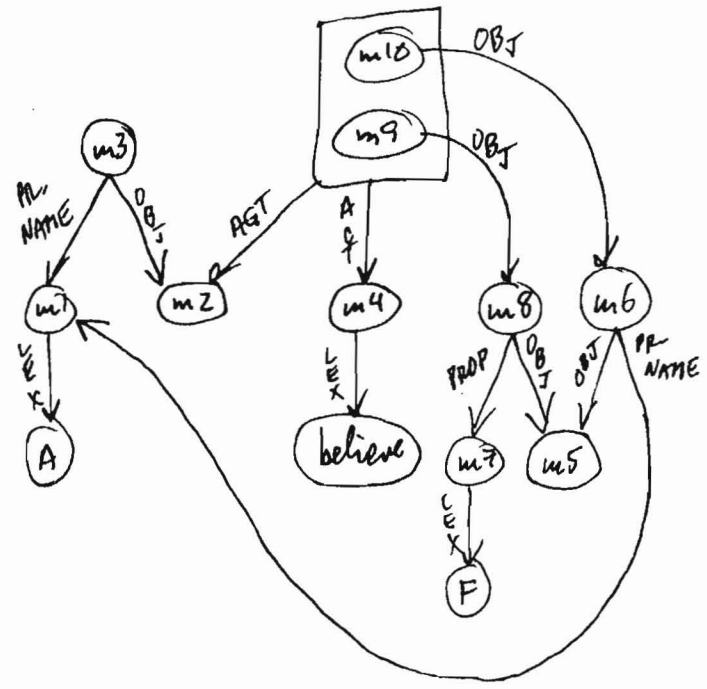

Fig. 11: A believes that A is F

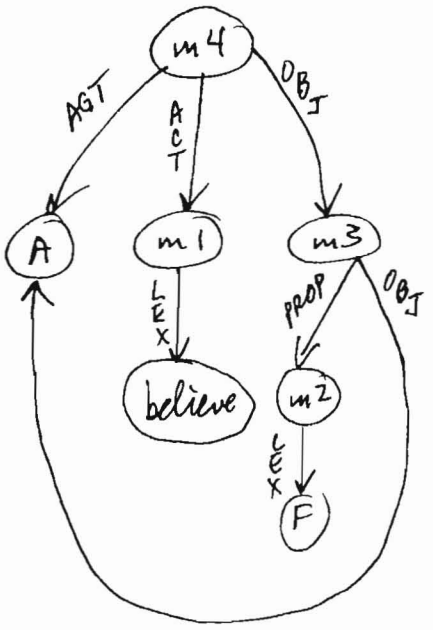

Fig. 12: A Maida & Shapiro network for 'A believes that * is F'

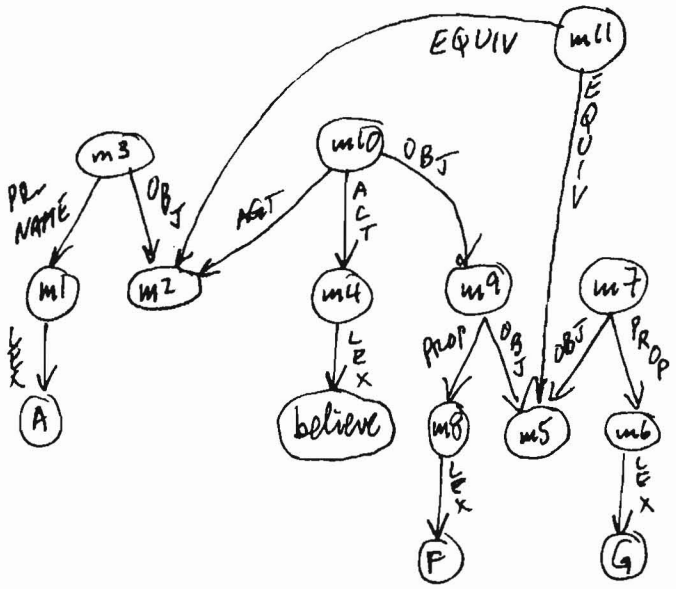

Fig. 13 Cassie's beliefs that:
A believes of someone who Cassie believes to be A (viz., someone who Cassie believes to have property G) that F

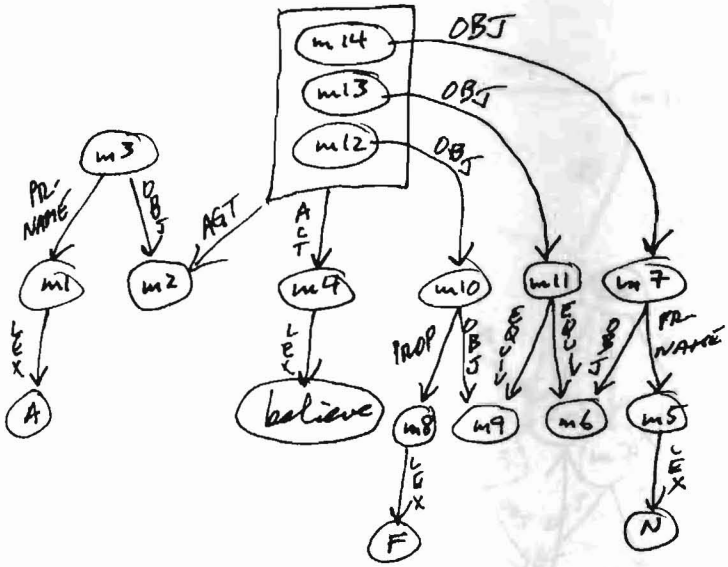

Fig. 14: A believes (de dicto) that something that is named 'N' is (the same as something that is) F

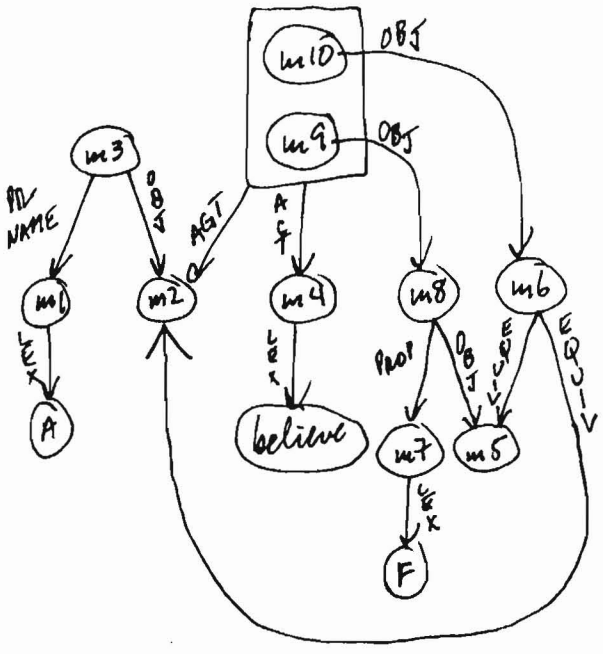

Fig. 15: A believes (de dicto) that something that is * is (the same as something that is) F

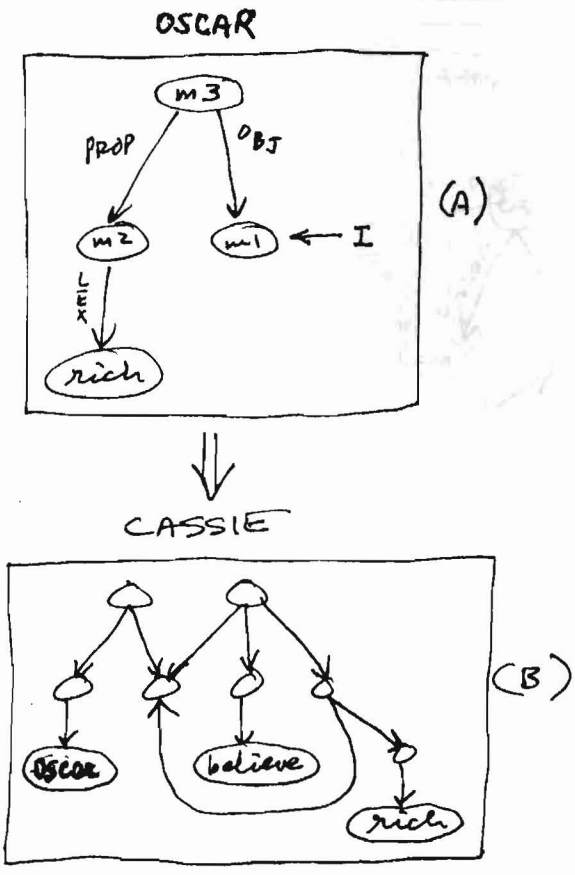

Fig. 16A: Oscar's belief: I am rich
B: Cassie's belief: Oscar believes that * is rich

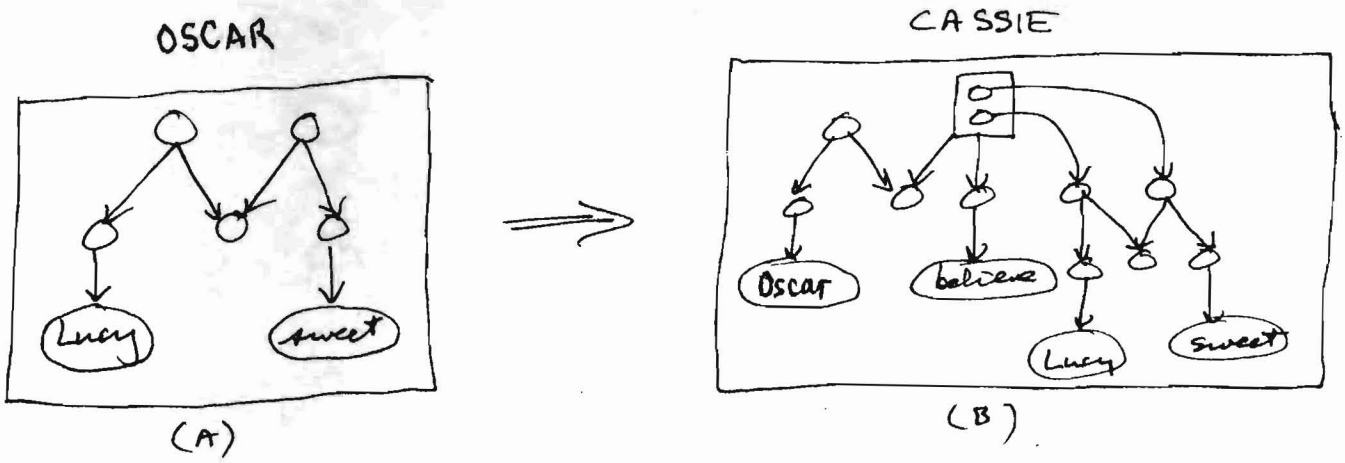

Fig. 17A: Oscar's belief: Lucy is sweet.

B: Cassie's belief: Oscar believes that Lucy is sweet.

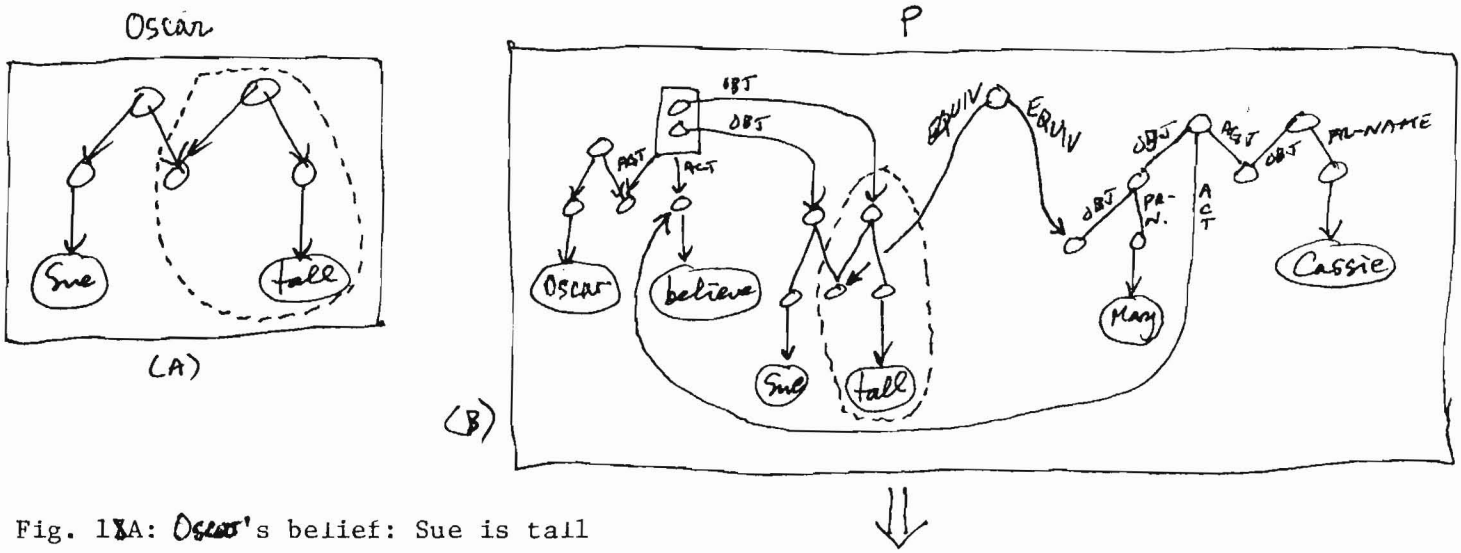

Fig. 18A: Oscar's belief: Sue is tall

B: P's beliefs: Oscar believes that Sue is tall and Oscar's Sue is Cassie's Mary

C: Cassie's belief: Oscar believes of Mary that she is tall

(Note: The networks in dashes correspond to each other)

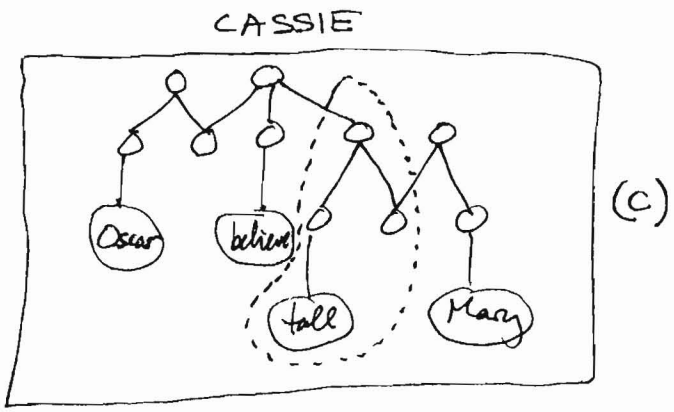

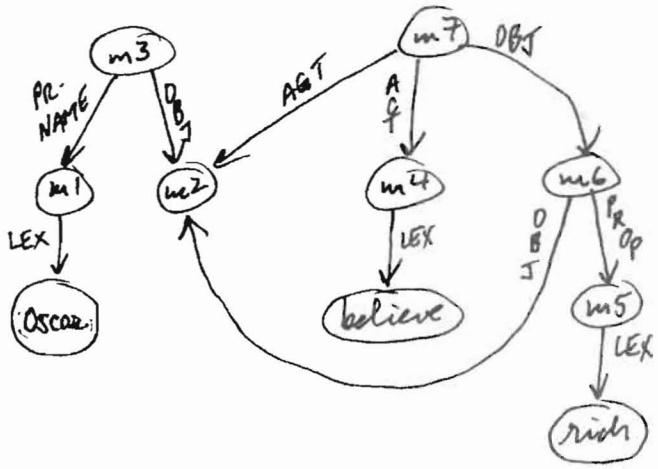

Fig. 19: Oscar believes of Oscar that he is rich

(?)

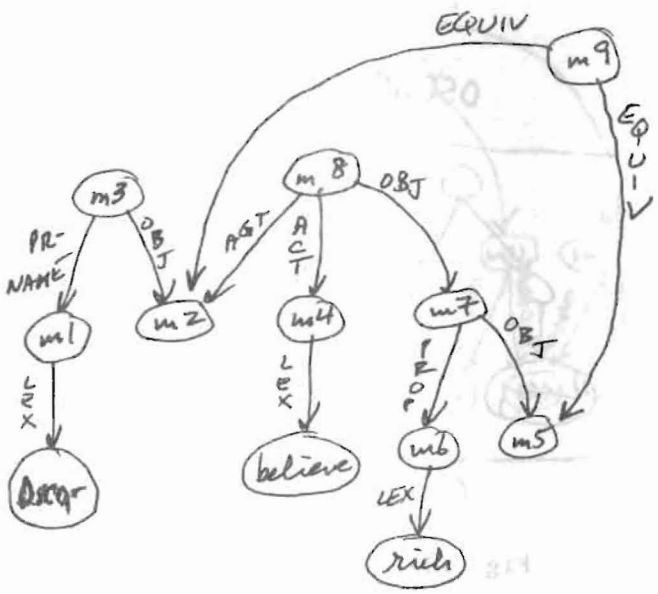

Fig. 20: Oscar believes of Oscar that he is rich

(?)

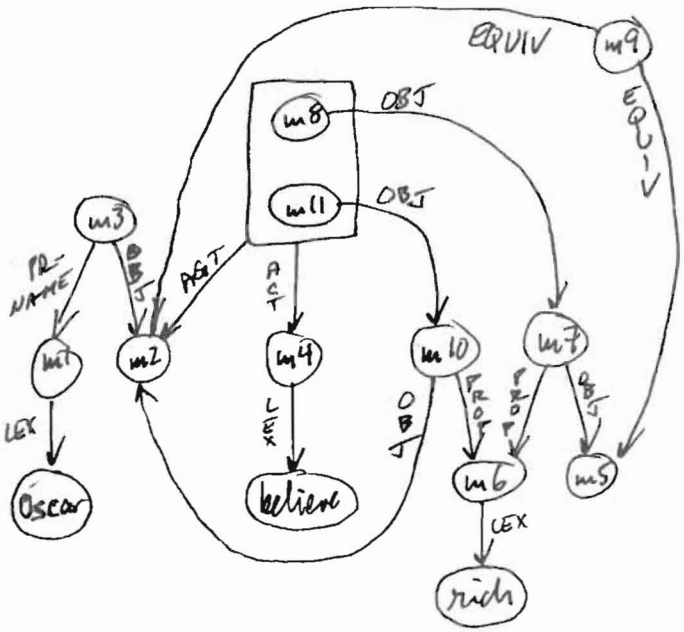

Fig. 21: Oscar believes of (the person who is) Oscar that he is rich (m8) and Oscar believes that he* is rich (m11)

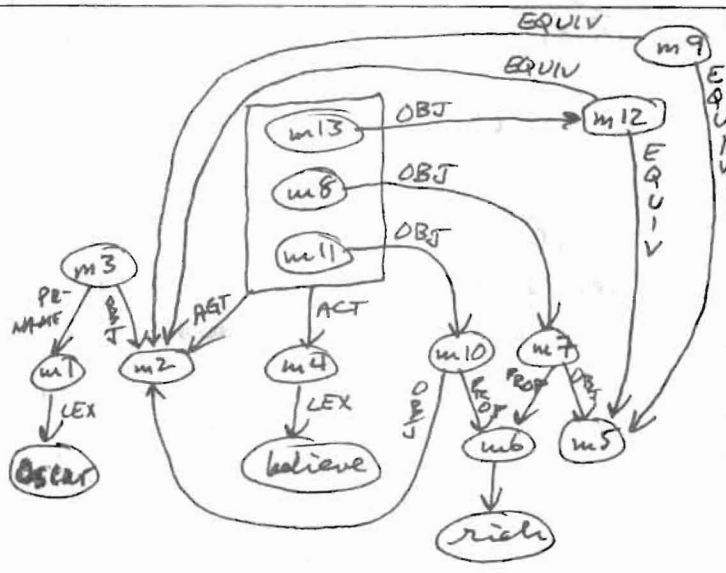

Fig. 22: Oscar believes of Oscar that he is rich (m8) and Oscar believes that he* is rich (m11) and Oscar believes that he* is the person whom he* believes to be rich (m13)