Representations of Intentions, Representations as Intentions, and Propositional Attitudes

by

J.A. Barnden
Department of Computer Science
Indiana University
Bloomington, IN 47405

TECHNICAL REPORT NO. 172

REPRESENTATIONS OF INTENSIONS, REPRESENTATIONS AS INTENSIONS, AND PROPOSITIONAL ATTITUDES

by

J.A. Barnden
June, 1985

REPRESENTATIONS OF INTENSIONS, REPRESENTATIONS AS INTENSIONS, and PROPOSITIONAL ATTITUDES

J.A. Barnden

Key phrases: beliefs, semantic networks, intensions, philosophy of representation.

ABSTRACT

Two views of declarative representation schemes have been important in theoretical psychology and artificial intelligence. In one the scheme contains expressions denoting cognitive structures in some cognitive agent. In the other the scheme's structures themselves act as cognitive structures in the agent, and can denote entities in the external world. We show that failure to distinguish carefully between these two views can lead to major difficulties in the representation of propositional attitudes (e.g. beliefs) held by cognitive agents. The phenomenon is presented by way of a criticism of the intensional semantic network scheme of Maida and Shapiro.

I INTRODUCTION

Declarative representations have been thought of in at least two distinct ways:

(R-OF):

as descriptions *of* cognitive structures (i.e. structures used by a cognitive agent in the course of cognitive processing);

(R-AS):

as cognitive structures.[1]

Theoretical psychology has made heavy use of declarative representations such as semantic networks. (See (Johnson-Laird et al. 1984) for a recent review.) Typically, one way we can view such networks is as R-OF representations. That is, we consider the networks to be formal notation whose semantics is based on the denotation of mental entities (e.g. concepts) by nodes or subnetworks. Thus we might have the "John" node in a network denoting the John concept in the mind. This concept in turn has as its extension the particular person John. It is also typically possible, however, to regard the networks as R-AS representations. In this case, certain entities in the brain/mind are taken to be *concrete realizations* (or *implementations*) of the networks. (E.g. in a quasi-neurophysiological account, on the lines of the schemes in (Hinton & Anderson 1981), abstract semantic networks are implemented as patterns of neural connectivity and/or activity.) Nodes in the abstract networks are taken to be *abstractions from* parts of those concrete realizations, rather than being taken to be formal objects that *denote* those parts according to some semantic function. An important consequence is that it is now proper to take the John node as *being* (part of) the John concept, and as *directly* denoting the person John.

To turn to declarative representations as used in AI, it is not surprising that there should be more of an emphasis on the R-AS view. Thus, for example, if one were to use Schubert's network scheme (Schubert et al 1979) in an AI program, the program would operate in part by manipulating pieces of network, and a denotational semantics of the network would be based on the denotation of entities in some "external world" by nodes in the networks. The same goes for a program based on network schemes such as that of (Hendrix 1979) or on logic representation schemes. However, some researchers present themselves as espousing the R-OF view. A major example of this is (Maida & Shapiro 1982). One of Maida and Shapiro's main tenets is that nodes in their networks denote intensions (concepts) as opposed to "extensions" (things in the "external world").[2] That is, the view taken would seem on this basis to be R-OF. However, we will see that Maida and Shapiro seem to be affected by the R-AS view as well, and that they are therefore led into grave difficulties.

There are other cases in which it is not clear that a representation researcher has, where appropriate, properly sorted out the distinction between the R-OF and R-AS views in his/her theorizing. (One or two examples are mentioned in section III.) It so happens that confusion between the two views, or even the maintenance of both views simultaneously, often does *not* cause difficulty in work on representation. With reference

---

[1] The distinction between R-AS and R-OF is related to the use/mention distinction for linguistic expressions.

[2] We will have cause later on to be worried about the meanings of the phrases "extension" and "external world".

to the depiction of the two views in Fig. 1, there is often no harm in occasionally identifying the (R-OF) representation scheme RS with the cognitive system CS it directly represents, thereby identifying the indirect-denotation link from RS to W with the direct-denotation link from RS to W in the R-AS case, and supressing the direct-denotation link from RS to CS.

However, as I show in this paper, great care must be taken with the distinction in a certain important area of representation theory: that of the representation of "propositional attitudes" (beliefs, desires, etc.). The reason the distinction sets traps in this area is that a given agent, A, can represent the propositional attitudes of other agents, and this representation can take the form of description (at some level of abstraction) of the other agents' cognitive structures. Under these conditions, if we are studying an R-AS representation scheme for A (i.e. the scheme is used by A), then this same scheme is R-OF for other agents. Worse, the scheme is R-OF for A as well as being R-AS, because A can reason about its own propositional attitudes. The complexity of the situation can lead us into confusion if we are careless about whether a representation scheme is R-AS or R-OF for given agents.

This paper deals with a particular sort of confusion that can arise. Suppose we are dealing with an R-OF representation scheme, (for a particular agent), but that (perhaps unconsciously) we treat part of the system as if it were R-AS, in the way explained above. What can happen is, as I will show, that a CS-to-W direct-denotation relationship encountered while we are under the spell of R-AS can be mistakenly "carried back" into an *RS-to-CS* direct-denotation relationship under the R-OF view, rather than to an indirect-denotation relationship under the R-OF view. This erroneous transfer can happen when the CS-to-W denotation is from a concept d in CS to a concept c in W, where, to make matters worse, c may be in CS too. We will see that the error has the effect of causing a whole level of conceptualization to be ignored.

The rest of the paper concentrates on detailing the phenomenon in the case of (Maida & Shapiro 1982). It should be observed that one of the main goals of (Maida & Shapiro 1982) is to devise a representation scheme that copes adequately with propositional attitudes.

In Fig. 1 we have a direct-denotation relationship between items in the R-AS representation scheme RS and entities in the world W. Now, part of our discussion will use the idea that such a representational item (e.g. a node in a network) is *part of a concept extending to the world entity.*[3] E.g. a node denoting John would be part of some entire concept the agent has of John. Another part of the concept might, for instance, be a visual image of his face or a procedure tailor-made for inferences about John.[4] We may then regard the denotation relationship from the node to the world entity as a part or aspect of the extending-to relationship from the concept to the entity. In fact, as a harmless over-simplification for the purposes of this paper, we will not distinguish between the node and the concept of which it is a part, and will thus conflate the denotation relationship and the extending-to relationship.

---

[3] When a concept c has an entity e as its extension, we say that c *extends to* e.

[4] I make no commitment to the nature of other parts. Also, it is in principle possible for the node to be the whole of the concept, so "part of" should be taken non-strictly.

II DISCUSSION OF THE MAIDA & SHAPIRO SYSTEM

A. The Main Difficulties

The proposition that John is taller than Mary would appear in a Maida & Shapiro network in the way shown in Fig. 2. Here the John, Mary and taller-than nodes denote the concept of John, the concept of Mary, and the concept of taller-than. The "head" node denotes the proposition. (A proposition is a sort of concept.) Under certain conditions on the embedding of this structure in the whole network, which we need not go into here, the truth value of the proposition would be taken to be "true". (In fact, a proposition has as its extension its truth value.) Fig. 3 shows the way the proposition that Bill believes that John is taller than Mary could appear. We reach our first difficulty at this point. It is to do with the semantics of this network substructure. Rather than dealing directly with a precise semantics, I will give an informal, simplified account. Consider the proposition denoted by the top node in Fig. 3. This proposition states a belief relationship between two entities. We ask: what sort of entities are they? One is a person and the other is a proposition. Notice here that the latter entity is just the concept denoted by the TJM node, whereas the former is the *extension* of the concept denoted by the Bill node. Thus, in determining what a proposition denoted by a node states, we sometimes "dereference" the concepts denoted by argument nodes and sometimes we do not.

This non-uniform dereferencing counts as a difficulty, because it forces cognitive processing mechanisms that act on the the mental entities denoted by the networks to be in some sense be aware of the need to dereference in some cases but not in others. An example of such a mechanism might be a system that translates cognitive structures (represented by Maida and Shapiro's networks) into natural-language statements. In the Fig. 3 example, we do not want the language generator coming out with a statement to the effect that a concept of Bill believes something or to the effect that Bill believes some truth-value (the extension of the TJM proposition)! However, things are not yet too bad, since we can simply take the view that processing mechanisms know that certain argument positions of certain relationships are of type "do-not-deref". Thus, there is some unwelcome complication, but not much.

But consider now the represention of the proposition that Mike's favourite proposition is more complex than Kevin's favourite proposition. This sort of example, where there are definite descriptions of propositions rather than explicit displays of them, is not considered in (Maida & Shapiro 1982), but Fig. 4 shows the network structure that would presumably be used. It is essential to realize here that MFP does *not* denote Mike's favourite proposition, but rather the concept of Mike's favourite proposition. (This is by analogy with what we would have if instead of "favoutite-proposition-of" we had "telephone-number-of": the node MFP would denote the concept of Mike's telephone number, according to (Maida & Shapiro 1982). Note also that the phrase "the concept of Mike's favourite proposition" should be taken as conveying that the concept is or includes a characterization of the proposition *as* Mike's favourite proposition.) Thus, MFP denotes a concept that extends to a concept (that itself extends to a truth value).[5] Notice that in saying what the proposition denoted by the top node in Fig. 4 is

[5] It would be possible to have MFP denoting Mike's favourite proposition directly, but on pain of taking the ad hoc step of making "favourite-proposition-of" act differently from

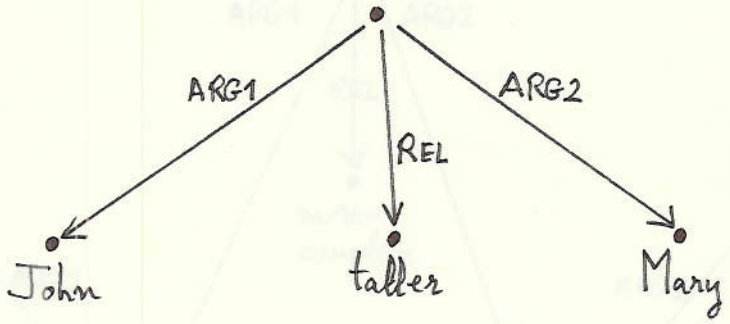

Fig. 2.

In this and later figures we simplify Maida & Shapiro's networks slightly (by abandoning LEX arcs).

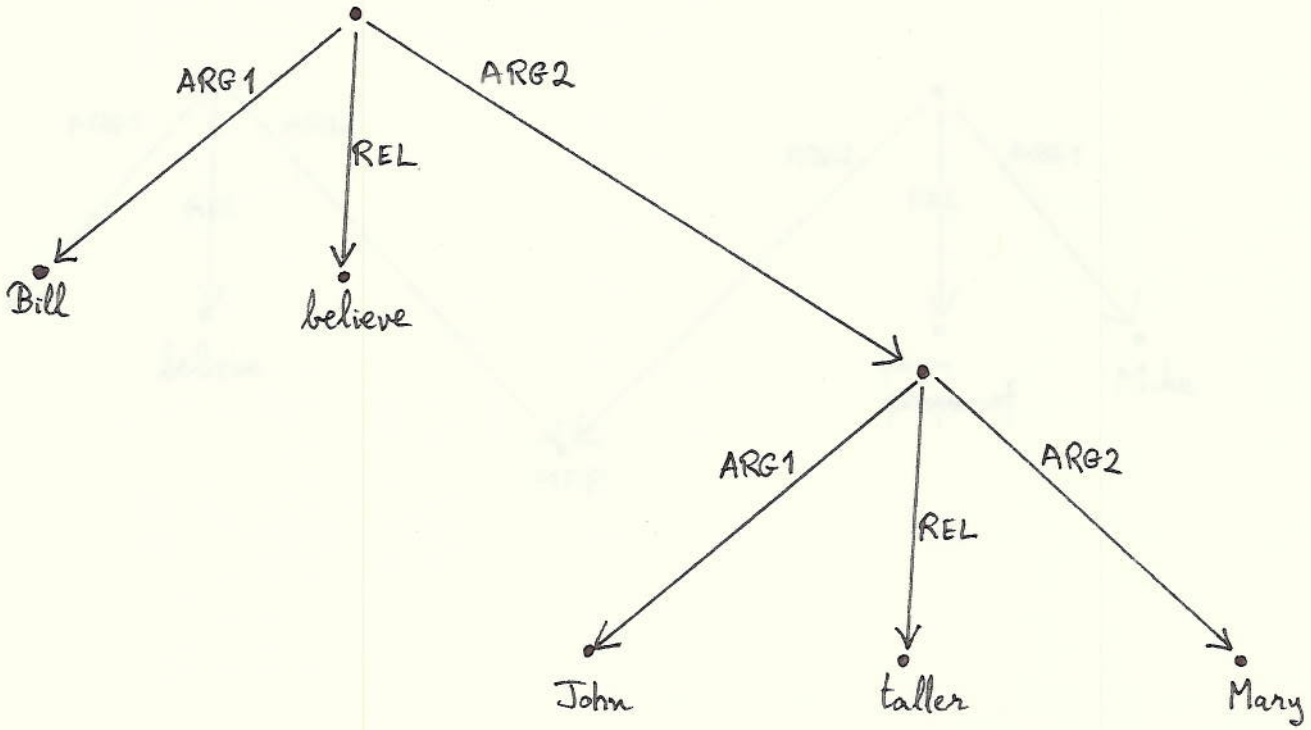

Fig. 3.

about, we must dereference MFP and KFP. The predicate "more-complex-than" is like the predicate "taller-than" in that it does not have do-not-deref argument-positions. Now the Fig. 4 example does not in itself cause difficulty; but what are we to make of the task of representing the proposition that Bill believes Mike's favourite proposition? Suppose we were to use the structure shown in Fig. 5. In saying what the proposition denoted by the top node states, we *do* now have to dereference the concept denoted by the ARG2 node in the belief structure, in contrast to the case of Fig. 3.[6] The simple suggestion of having do-not-deref argument positions is thus inadequate. An alternative technique that would cope with the present problem, as well as with the Fig. 3 difficulty, is to refrain from dereferencing concepts denoted by just those nodes that send out REL arcs. We may allow for the sake of argument that this interpretation technique would result in a trouble-free system (a point which would need careful examination) – our claim is not that the system cannot be given a rational interpretation, but rather that there are strong grounds for suspicion that the representation scheme is ill-conceived in part.

But we now get into trouble with the proposition that Kevin's favourite proposition is that John is taller than Mary. Neither the structure in Fig. 6(a) nor that in Fig. 6(b) are satisfactory. The problem in Fig. 6(a) is that, by analogy with Fig. 2, TJM denotes the proposition that John is taller than Mary, but equally, by analogy with the Fig. 4 example, it should denote a concept of some proposition. The problem in Fig. 6(b) is that, according to the account given by Maida and Shapiro of their EQUIV feature, the structure states that the extension of the proposition that John is taller than Mary is the same as the extension of the concept denoted by node KFP. But this would state a nonsense: that a truth-value is Kevin's favourite proposition. We could consider getting round this problem by altering the way that EQUIV statements are interpreted and using the REL fix mentioned a moment ago, so that in Fig. 6(b) the EQUIV structure would state that the proposition that John is taller than Mary is *itself* the same thing as the extension of the concept denoted by node KFP. However, this just introduces further ad-hoc-ness.

One way that might be suggested for getting out of the difficulty is to introduce an "extension-of" relation that could be used to express explicit dereferencing. Using this, we would render the proposition that Bill believes Mike's favourite proposition by the structure shown in Fig. 7. (Note that extension-of dereferences on neither argument.) Extension-of could also be used in the representation of the proposition that Kevin's favourite proposition is that John is taller than Mary. The trouble with the suggestion, of course, is that, for consistency, we should be able to have structures like the one in Fig. 8, where the J node now denotes the *person* John: but this contravenes the strongly held tenet that nodes should not denote things like people, but only concepts of them! We can wiggle out of this by allowing extension-of to have as ARG2 nodes only nodes

---

"telephone-number-of". Also, it is not difficult to see that Maida and Shapiro's "Uniqueness Principle" would be violated under some situations in which Mike's favourite proposition is the same as Kevin's.

[6] The proposition B denoted by the top node in Fig. 5 states that Bill believes that proposition P that is Kevin's favourite. The characterization of P as Kevin's favourite is thus in some sense part of B; nevertheless, it is P itself that is stated to be believed, not the characterization. It is in this sense that in Fig. 5 we have to dereference the concept denoted by the ARG2 node of the belief structure; but this should not be taken to mean that the characterization involved in that concept is entirely discarded.

denoting second-or-higher level concepts (concepts of concepts (of ...)). But there appears to be no independent argument to justify this restriction. A further difficulty is that, with this restriction, the extension-of relationship would be of no help in stating that the a particular truth-value is the extension of the proposition denoted by a node N sending a REL arc. This is because N would be the ARG2 node for the extension-of relationship, but denotes a first-level concept. Now (Maida & Shapiro 1982) does postulate a truth-value relation that is used to state what truth-values proposition have; but to use this relation as well as extension-of would put us in the messy position of expression some extending-to relationships in one way and others in another way.

What these considerations highlight very clearly is an over-simple attitude towards "intensions", "extensions" and the "world". The philosophy of Maida and Shapiro fails to take into account the fact that some intensions can have as their extensions other intensions, and that therefore the world W in Fig. 1 contains not only the things, like people, that Maida and Shapiro count as extensions, but also concepts, like the propositions in the examples we have discussed (so W and CS intersect).

B. The Connection with R-OF and R-AS

There is a way of modifying the Maida and Shapiro system to get over the difficulties presented: namely, to introduce an extra level of intension under every proposition node.[7] That is, we make the head node of any proposition structure denote not the proposition itself but rather a concept of that proposition. For example, in Fig. 2 the top node would now denote a *concept of* the proposition formed by applying the taller-than predicate to John and Mary. Then, in Fig. 3 both argument positions of the belief structure must be dereferenced in order to discern what the proposition states. The problems with favourite-proposition-of also go away.

The extra elaborateness is a natural consequence of having an R-OF representation scheme, and the fact that the mental entities themselves are part of the world represented by the mind and therefore subject to being intensionally described in different ways, just as ordinary objects can. This latter point can be seen operating in the fact that the top node of Fig.2 is now to be regarded as a definite-description node, on a par with the MFP node in Fig. 4. MFP denotes the concept of the favourite proposition of Mike – *characterized as such* by the concept. The top node in Fig. 2 denotes the concept of the proposition formed from the taller-than relationship and persons John and Mary – the proposition being characterized as such by the concept.

We conjecture that Maida and Shapiro devised much of their system with the R-OF view firmly in mind, but were somewhat affected by the R-AS view when considering the denotations of proposition nodes. They correctly realized that a representation scheme able to cope with the representation of the propositional attitudes of various agents can be achieved by having structures denoting intensions explicitly. (The same basic insight is used also in other systems, such as that of (Creary 1979), although in a different way in detail.) Now, in an R-AS representation scheme one consequence of this realization is to have nodes (or other appropriate items) denoting propositions (which are

[7] The modification is not purported to result in a trouble-free scheme. It is not the purpose of this paper to detail a particular representation scheme.

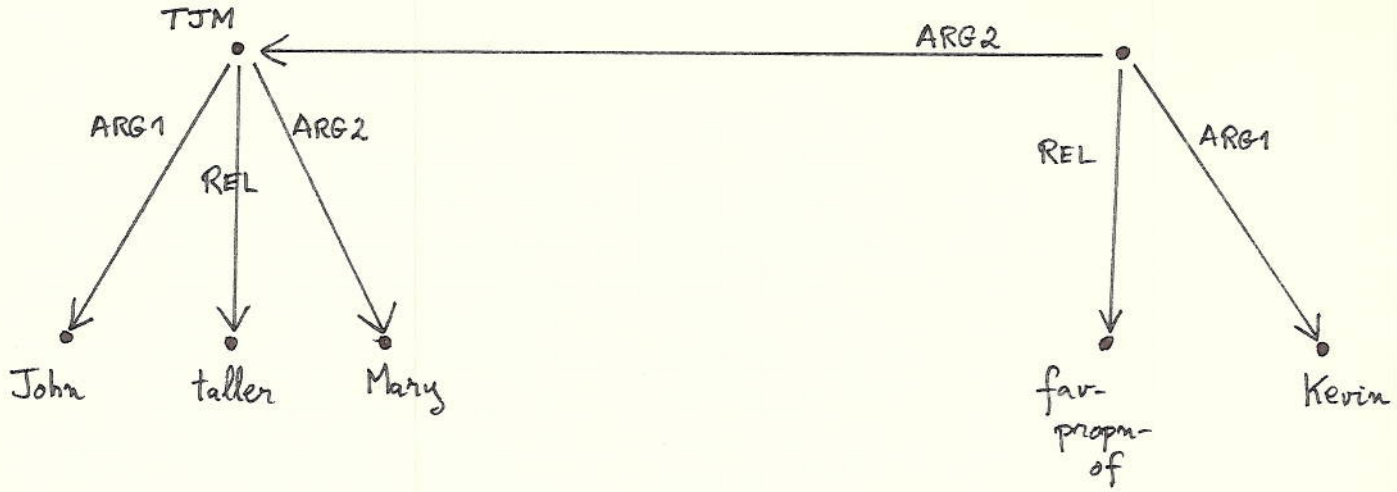
TJM
ARG2
ARG1
REL
ARG2
REL
ARG1
John
taller
Mary
fav-
propn-
of
Kevin

Fig. 6(a).

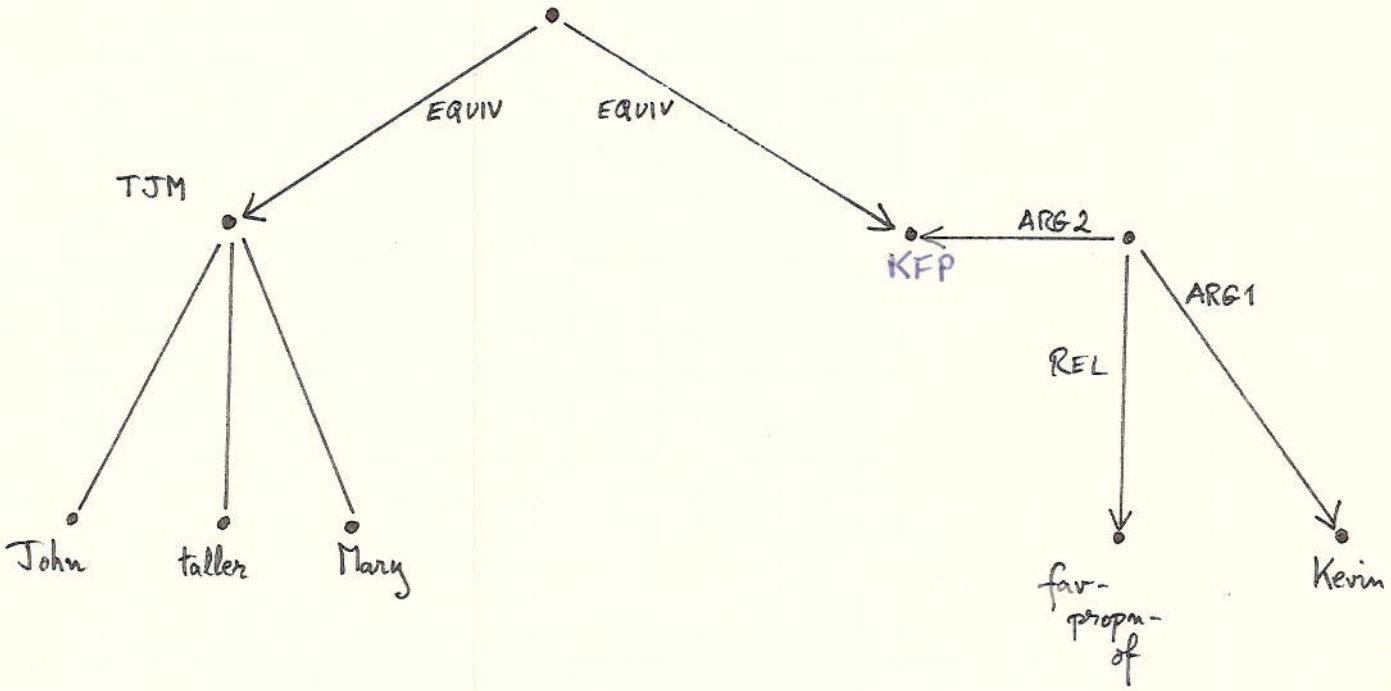
EQUIV
EQUIV
TJM
ARG2
KFP
ARG1
REL
John
taller
Mary
fav-
propn-
of
Kevin

Fig. 6(b).

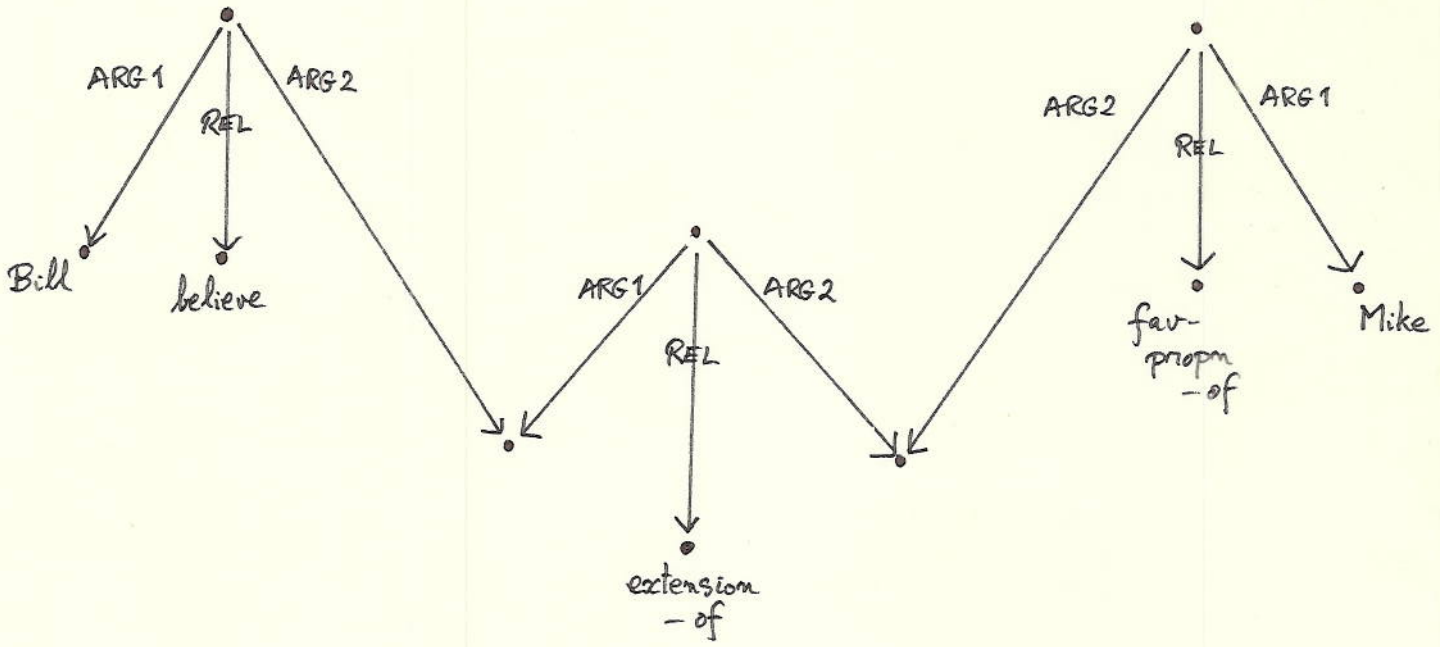
ARG1
REL
ARG2
Bill
believe
ARG1
REL
ARG2
extension
-of
ARG2
REL
ARG1
fav-
prop
-of
Mike

Fig. 7.

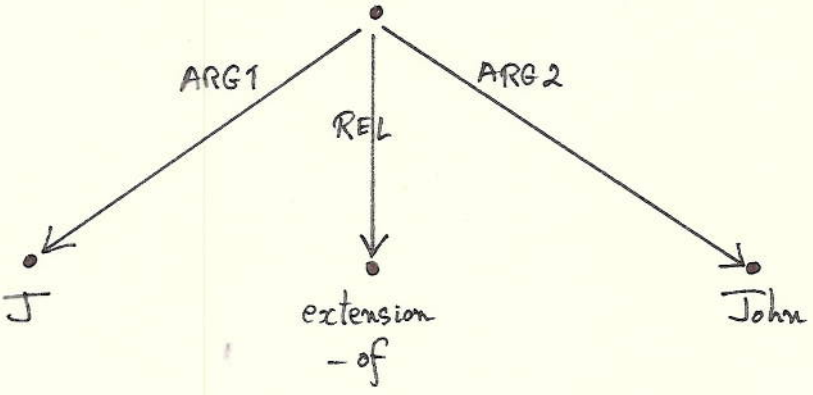
ARG1
REL
ARG2
J
extension
-of
John

Fig. 8.

a type of intension). For instance, the scheme could represent Bill's believing that John is taller than Mary by means of a structure in which there is a node (say) denoting Bill and a node denoting the proposition c that John is taller than Mary. But, by the comments at the end of section A, we regard the latter node as (part of) a concept d extending to c. In moving to an R-OF view, we need a node dd in the RS to directly denote d, *not* the proposition c extended to by that concept. In the R-OF scheme we still have the denotation relationship from d to c, and merely an indirect-denotation relationship from dd to c. What has happened is that the d-to-c denotation in the R-AS view has been "carried back" to yield, in the R-OF view, both the d-to-c denotation itself and the indirect dd-to-c denotation. (See Fig. 9.) However, we conjecture that what Maida and Shapiro have done is, in an unconscious move from R-AS to R-OF, to "carry back" the d-c direct-denotation into an erroneous dd-to-c direct-denotation, making c usurp the role of d as denotatum of dd, and losing d entirely, as shown in Fig. 9. (Thus, in our particular example, Maida and Shapiro have a node dd denoting the proposition c that John is taller than Mary, whereas according to the modification suggested at the head of this section, we should have a node dd denoting a concept d denoting that proposition.) The erroneous step was facilitated by the fact that d and c are confusingly similar in that they are both intensions that have something to do with propositions, and both can be taken to be concepts within the agent being modelled.

We have seen that Maida and Shapiro purport, when talking theoretically, to suscribe to the R-OF view of their scheme, while, we conjecture, suffering ill-effects from "contamination" by a partial R-AS view. There is independent circumstantial evidence for this contamination. For, the whole tenor of the work indicates that Maida and Shapiro have it in mind for AI programs to *use* their representation scheme. Also, on p.300/301 of (Maida & Shapiro 1982) we read about a hypothetical robot in which there are connections between network nodes and sensors and effectors. On p.319 of the paper, nodes are talked about as if they are *in* cognitive agents rather than just being items in a theory *about* the cognitive agents. To give the authors the benefit of the doubt we could say that they are, merely, loosely talking about nodes when what they really mean to talk about is the mental entities denoted by those nodes. If so, we might have expected an explicit caveat to this effect.

Finally, it should be clear that, although the relation favourite-proposition-of does not itself have much practical significance, the problems it subtends will crop up with other relations allowing descriptions of propositions.

III OTHER WORK

(Creary 1979) presents an (almost) first-order logic scheme where terms can denote concepts (including propositions). His scheme is able to make all the distinctions Maida & Shapiro's is capable of making. His scheme escapes the sort of difficulty discussed in this paper (although not others, as shown in (Barnden 1983)) partly as a result of sticking firmly to the R-AS view. By sticking to the R-AS view, he avoids the extra level of representation involved in the R-OF view. Similarly, one's theorizing about the FOL system (Weyrauch 1980) – in so far as it is applied to propositional attitude representation – and about the belief system of (Konolige 1983) is simplified by virtue of the

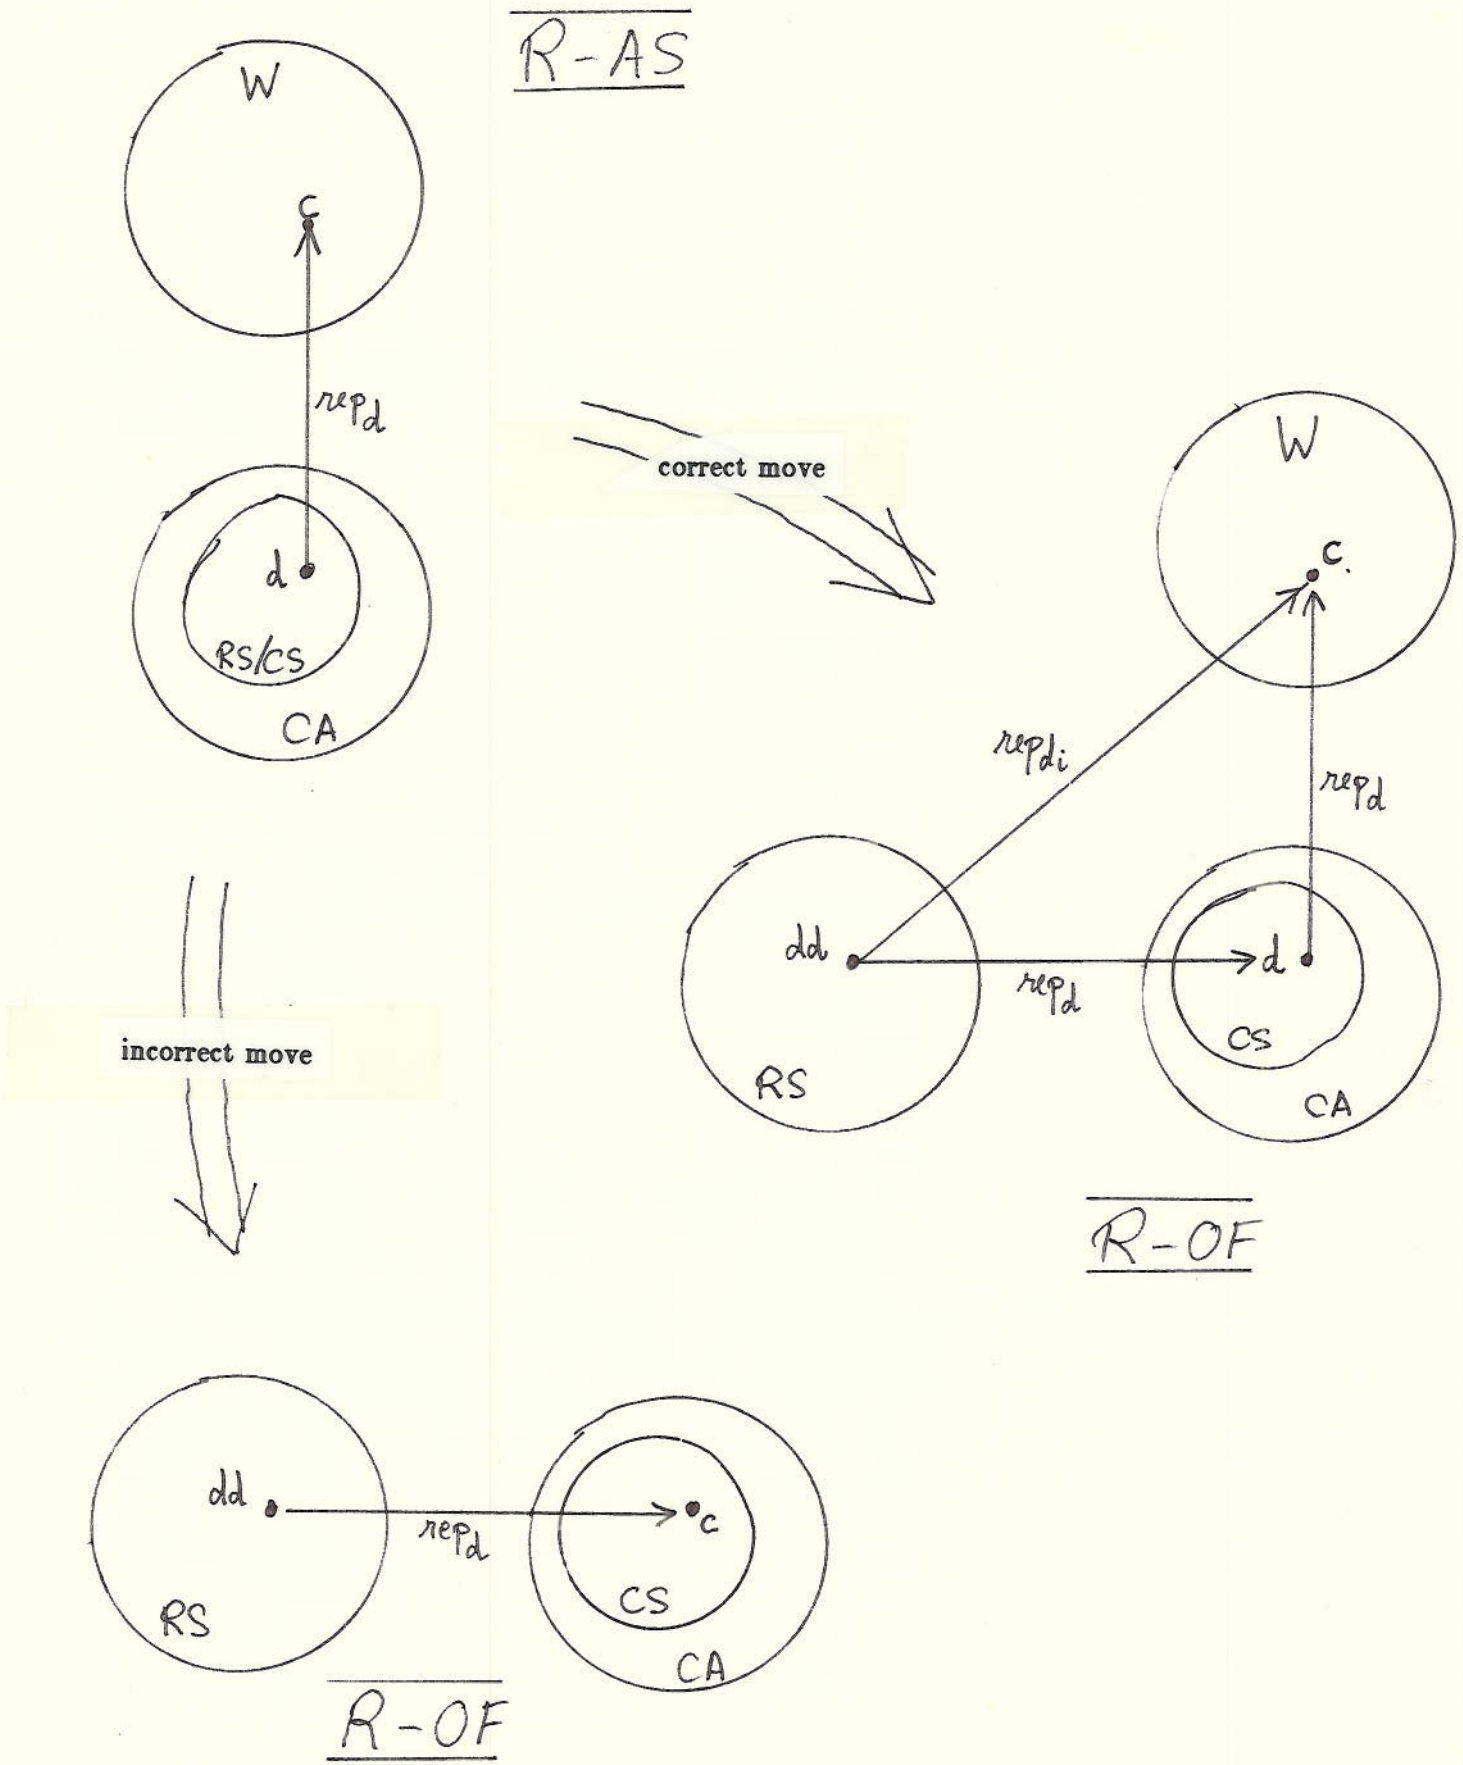
R-AS
W
c
rep_d
d
RS/CS
CA
correct move
W
c.
rep_di
rep_d
dd
rep_d
d
RS
CS
CA
R-OF
incorrect move
dd
rep_d
c
RS
CS
CA
R-OF

Fig. 9.

representation schemes being R-AS.

Brachman's KL-ONE scheme is not focused on propositional attitudes in the way Maida and Shapiro's is. However, in its sophisticated treament of conceptual structures it is a good candidate for application to the subtle and complex issues raised by propositional attitudes. It is therefore with some alarm that in a paper such as (Brachman 1979) we find, at best, severe unclarity of presentation. On pp.34/5 of (Brachman 1979) we are told in one breath that "Concepts" (which, note, are formal objects in the representation scheme) represent intensional objects, not "extensional (world) objects", while in another we are told that Concepts represent objects, attributes and relationships of the the domain being modelled; also, we are told that, say, the ARC-DE-TRIOMPHE Individual Concept "denotes" the real Arc de Triomphe. Analysis of the text is difficult, because it may be that Brachman takes "denote" to mean something different from "represent", and may be using more than one meaning for "represent" alone (or may simply be using the word carelessly). However, since KL-ONE seems to have been partly intended for use by AI programs, one is justified in suspecting that Brachman is (in the cited paper) trying simultaneously to hold an R-OF view and an R-AS view. This feeling is reinforced by the use of the word "Concept" as a name for something that is claimed to *represent* (denote?) a concept (or "intensional object").

We have reason for being uncomfortable with the possible-world approach to knowledge (as a propositional attitude) presented in (Moore 1977). Quite apart from the traditional objections to possible-world approaches (e.g. that there are grave problems of implausible "omniscience" on the part of the cognitive agents – see e.g. (Linsky 1983)) we have the following consideration. The scheme presented contains terms that denote possible worlds, and is claimed to have certain practical, inferential advantages. Presumably this claim is support for the idea of using the scheme as a basis for an AI program. Suppose we were to construct such a program, and called it cognitive agent A. A is able to reason *explicitly* about possible worlds, and is surely to be taken to have explicit knowledge about possible worlds. However, the cognitive agents represented in the scheme do not reason about possible worlds. What, then, are we to make of the problem of representing an agent that knows something about what A itself knows? Are we to pretend that as far as all non-A agents are concerned, A knows nothing about possible worlds, even though A knows things about possible worlds? If so, why? We conjecture that Moore's representation scheme was conceived primarily as a theoretical tool for representing mental states (and actions) – and is therefore to be considered an R-OF scheme – and the consequences of an R-AS view of the scheme were not fully worked out (even though some of the claimed advantages are partly on the practical side, as opposed to being concerned with such matters as theoretical economy).

IV CONCLUSIONS and FUTURE DIRECTIONS

We observed that Maida and Shapiro's representation scheme has semantic irregularities, that must, moreover, have unfortunate practical consequences for processing mechanisms. The difficulties arose in cases of representing agents' propositional attitudes, and result from a combination of factors: over-simplification on the subject of intensions and extensions, and a failure to realize the full consequences of adopting an

R-OF view as opposed to an R-AS view.

My own current view about how propositional attitudes should be dealt with in an R-AS representation scheme is based on the (old) idea of expressions that denote expressions (e.g. by quoting them), the latter being expressions in the representation scheme itself. Thus, one way of stating that Bill believes that John is taller than Mary would be to have a representational structure (expression) in which one item is a direct quotation of some expression that could be used to state that John is taller than Mary. (See (Quine 1981) and (Perlis 1985) for proposals along these lines.) There are numerous technical and conceptual issues that arise in the development of this idea. For instance, there is the danger of introducing semantic paradoxes if certain types of predicate (e.g. a truth predicate) are allowed in the system. This particular issue might be dealt with by means of the semantic scheme proposed in, say, (Kripke 1975). Our concern in the present paper is not with such difficulties, however, but with the way some of our considerations here provide partial support for the denotation-of-expressions view.

Consider then that at the end of section I we claimed that a node (or other representational expression) in an R-AS scheme is to be regarded as part of a concept that the cognitive agent has of whatever entity is denoted by that expression. In saying this we are taking a concept to be some structure or other in the total cognitive system of the agent. But, we have all along taken a concept to be an intension, so we have implicitly adopted the view that an intension is a cognitive structure (as opposed, say, to some abstruse mathematical object such as a mapping from situations or possible worlds into truth values, or such as some equivalence class of expressions). Now, we could take the stance that this view of intensions was an over-simplification adopted for the informal purposes of this paper, and that a more precise expression of the considerations would involve a different view of intensions. However, I currently prefer the stance that intensions really should be taken to be cognitive structures (i.e. abstractions, at some high level, of the physical states of the cognitive agent). A consequence of this is that we will not necessarily be able to say that two different agents both hold an intension – but intensions in different agents may be more or less similar. How, then, should the representation scheme in the agent represent propositional attitudes of agents? Following the line taken in (Creary 1979) (arguably the most advanced and promising AI approach so far to the problems of propositional attitudes), we assert that the representation should include terms denoting intensions. There remains the question of which agents should own these intensions. If, for instance, Mike's scheme contains a statement that Bill believes that John is taller than Mary, we want Mike to contain a representational term denoting some cognitive structure (intension) C stating that John is taller than Mary. The natural claim is that C should be owned by Bill. (Hence Mike's representation scheme is R-OF for Bill, as it contains an item denoting a mental entity within Bill.) However, Mike presumably does not have detailed knowledge of Bill's intensions, so that C would have to be imperfectly characterized by some sort of description in Mike. An alternative is that C should belong to Mike himself, the idea being that Bill owns an intension *somewhat like* C, where Mike *makes no specific commitment about the nature of this likeness*. Under this proposal, we do not even in principle pretend that agents have exact knowledge of each other's propositional attitudes, and we now have a scheme that is only approximately R-OF for Bill. Finally, if we let C belong to Mike, then C could possibly be denoted by being directly quoted (under some notion of "quote" for expressions more general than the symbol strings we usually apply quotation to), although there may be other useful ways of denoting C (e.g. by use of constructor

functions).

In saying an intension is a cognitive structure, I do mean to imply that is is a cognitive structure in full goriness. That is, it may include such things as perceptual images (whatever *they* are), "attached" inference procedures, motor procedures, ... There is therefore a major set of problems to be resolved in detailing how such objects can be denoted and, in particular, quoted. As a methodological strategy for avoiding these problems temporarily while allowing others (e.g. semantic paradoxes) to be attended to, I propose that we should at present pretend that cognitive structures consist merely of representational structures in some logic/net/frame/... representation scheme. The structures denoted and quoted are to be regarded as intensions, as long as it is realized that this is just a temporary, over-simplified view.

REFERENCES

Barnden, J.A. Intensions as such: an outline. In *Procs. 8th. Int. Joint Conf. on Artificial Intelligence,* Karlsruhe, 1983.

Brachman, R.J. On the epistemological status of semantic networks. In (Findler 1979).

Creary, L.G. Propositional attitudes: Fregean representation and simulative reasoning. In *Procs. 6th. Int. Joint Conf. on Artificial Intelligence,* Tokyo, 1979.

Findler, N.V. (ed.). *Associative Networks.* Academic Press: New York, 1979.

Hendrix, G.G. Encoding knowledge in partitioned networks. In (Findler 1979).

Hinton, G.E. & Anderson. J.A. (eds) *Parallel Models of Associative Memory.* Lawrence Erlbaum: Hillsdale, NJ, 1981.

Johnson-Laird, P.N., Herrman, D.J. & Chaffin, R. Only connections: a critique of semantic networks. *Psych. Bull. 96* (2), 292-315, 1984.

Konolige, K. A deductive model of belief. In *Procs. 8th. Int. Joint Conf. on Artificial Intelligence,* Karlsruhe, 1983.

Kripke, S.A. Outline of a theory of truth. *J. Philosophy, 72,* 217-21, 1975.

Linsky, L. *Oblique Contexts.* U. Chicago Press, 1983.

Maida, A.S. & Shapiro, S.C. Intensional concepts in propositional semantic networks. *Cognitive Science 6,* (4), 291-330, 1982.

Moore, R.C. Reasoning about knowledge and action. In *Procs. 5th. Int. Joint Conf. on Artificial Intelligence,* MIT, 1977.

Perlis, D. Languages with self-reference I: Foundations. *Artificial Intelligence, 25,* 301–322, 1985.

Quine, W.V. Intensions revisited. In W.V. Quine, *Theories and Things.* Harvard U. Press: Cambridge, Mass, 1981.

Schubert, L.K., Goebel, R.G. & Cercone, N.J. The structure and organization of a semantic network for comprehension and inference. In (Findler 1979).

Weyrauch, R.W. Prolegomena to a theory of mechanized formal reasoning. *Artificial Intelligence 13* (1), 133–170, 1980.