A VIEWPOINT DISTINCTION IN THE REPRESENTATION OF PROPOSITIONAL ATTITUDES

John A. Barnden

Computer Science Department
Indiana University
Bloomington, Indiana 47405

ABSTRACT

A representation scheme can be *used by* a cognitive agent as a basis for its normal, inbuilt cognitive processes. Also, a representation scheme can serve as a means for *describing* cognitive agents, in particular their "mental" states. A scheme can serve this second function either when it is itself naturally used by a cognitive agent (that reasons about agents), or when it is merely an artificial, theoretical tool used by a researcher. In designing a representation scheme one must pay very careful attention to two related questions: the question of whether, for any given agent, the scheme is used by the agent or is used to describe the agent (or both); and the question of whether the scheme is being used as a theoretical tool (as well as, perhaps, being used by agents). I show by example that representational pitfalls can be encountered when these questions are not clearly addressed. The examples revolve around Creary's logic-based scheme and Maida and Shapiro's semantic network scheme, both of which were designed primarily to facilitate the representation of propositional attitudes (beliefs, hopes, desires, etc.). However, the general points have wider application to schemes for propositional attitude representation. By appeal mainly to the Maida and Shapiro case I demonstrate also that it is possible to be misled by the ambiguity of whether "to represent" means "to denote" or "to be an ambassador/representative/abstraction of".

I TWO REPRESENTATIONAL VIEWPOINTS

For the sake of discussion I will view a (declarative) representation scheme as consisting of (a) a specification of a set of possible expressions together with (b) a specification of how expressions either denote (refer to) entities in some world or make assertions about such entities, and (c) rules about how expressions can be manipulated. Item (b) is the "semantics" of the scheme.

It is typical for the semantics of an AI representation scheme to make denoting expressions denote "ordinary" entities such as people, blocks, places, telephone-numbers, etc., and to take the assertions made by assertional expressions to be about such things. The scheme is (typically) viewed as instantiated in one or more cognitive agents that *use* the scheme in order to achieve their goals. This means in part that the scheme's expressions are abstractions from the internal, "mental" nature of each such cognitive agent X. A given expression, which might denote the person John for example, is an abstraction from a something-or-other internal to X.

Now let us assume, just for definiteness, that this something or other is X's concept of John, or at least part of it. Then this concept (or "intension") is deemed to have the person John as its "extension". Notice the three things we have here: the representational expression, X's concept of John, and the person John. AI researchers are accustomed to saying that the expression "represents" John, where the notion of representing is that of "denoting" — in the sense in which a term in a logic scheme denotes an entity according to some interpretation. However, it is also possible to consistently maintain that the expression "represents" X's John-*concept*, provided we realize that the notion of representation here is akin to "being a representative or ambassador of", "being a theoretical handle on" or "being an abstraction of" rather than to the notion of *denoting* (or referring to).

Of course, it is possible for a representation scheme's expressions to *denote* concepts, or, more generally, entities within or aspects of an agent Y's mental make-up. (Concomitantly, some expressions would make assertions about such things.) Important examples of schemes with this property are the concept-denoting schemes of McCarthy (1979) and Creary (1979). This sort of representation of concepts by expressions must be carefully distinguished from the ambassadorial/abstractional sort. Indeed, the scheme might be being *used* by an agent X, so that the scheme's concept-denoting expressions ambassadorially represent concepts (or other mental entities) in X but denote concepts in Y. (Extra difficulty arises when Y is actually X; but the two types of representation must still be distinguished.) Equally, the scheme could be in use by researchers as a *theoretical tool* for describing a domain that includes agent Y, and *not* used by any agent as part of its normal processing. Some expressions in the scheme then denote concepts, but none ambassadorially represent any.

The notion of ambassadorial representation applies also to assertional expressions in a scheme instantiated in an agent X. Such an expression abstracts from some assertional mental something-or-other in X. The expression does not of course usually assert anything *about* that something-or-other.

Part of the purpose of the paper is to demonstrate the sort of problems that can arise if the distinction between denotation and ambassadorial representation is not properly maintained (although of course this distinction is not new in itself). But another major, related goal is to show the importance of distinguishing between the following two possible views of a representation scheme:

(A) as something that is *used by* an agent as a basis for its normal cognitive processing;

(B) as a means for *describing* (making assertions about and/or denoting aspects of) agents, especially in the case when the scheme is used as an artificial, theoretical tool by a researcher.

On the way, I will demonstrate some representational infelicites that do not seem to have been noticed before.

The main discussion will centre on attempts to get representation schemes to embody information about "propositional attitudes" (beliefs, desires, hopes, etc.). This should be no surprise, in view of our concern with description of agents. The study can be seen as a contribution to our appreciation of the complex subtleties inherent in the topic of propositional attitudes. The points I will make are closely connected with observations made in the past by other people, but there will not be space to trace the connections. See (Barnden, 1986) for further discussion.

My main examples of difficulties arising in representation schemes come from the work of Maida & Shapiro (1982) and of Creary (1979). Before going on to this, it is worth pointing out one or two preliminary examples of how the ambiguity of "to represent" has appeared in the literature. A first instance is Brachman's writings on KL-ONE (Brachman, 1979; Brachman & Schmolze, 1985). The KL-ONE scheme is not focused on propositional attitudes. However, in its sophisticated treatment of conceptual structures it shows promise of being a good candidate for application to the subtle and complex issues raised by propositional attitudes, and Brachman & Schmolze (1985) state that the context mechanism should prove useful for propositional attitude representation. 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, that 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. 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) unclear as to whether his representational expressions denote concepts or ambassadorially represent them. This feeling is reinforced by the use of the word "Concept" as a name for a representational object that is claimed to *represent* (denote?), rather than to *be* (an abstraction of), a concept or "intensional object". The recent description of KL-ONE in Brachman & Schmolze (1985) also uses the term "represent" ambiguously in places.

The noted ambiguity is analogous to one that creeps into Johnson-Laird's work on mental models (Johnson-Laird, 1983). His theory is based on the idea that our minds manipulate both "propositional representations" (natural-language-sentence-like structures, which denote and assert in familiar ways) and structures called "mental models". Johnson-Laird talks much of the time as if mental models are models of proposition representations — using the term "models" in something akin to the model-theoretic sense of logic. The components of models are thus internal objects that are, presumably, to be considered as denoted by terms in the propositional representations. At the same time, Johnson-Laird often talks of mental models as "representing" states of affairs in the outside world. [For just one example, see (Johnson-Laird, 1983:p.419).] Now, it seems most consistent with his general viewpoint to take him to mean that mental models are abstractions from outside-world states of affairs, and thus to "ambassadorially" represent them in a sense; rather than that mental models denote, and make assertions about, outside-world entities. However, this view is somewhat at odds with his (more tentative) suggestions that mental models could contain propositional-style elements, such as numerical tokens. The point is that such a suggestion seems to imply that a model can be partially denotative, as well as ambassadorial. This mixing may well be what is needed — but the point is that the issues should be made clear, and the various notions of representation in use should be properly distinguished.

II THE MAIDA AND SHAPIRO SYSTEM

Here I demonstrate some difficulties in the semantic net scheme of (Maida & Shapiro 1982). [See also Rapaport & Shapiro (1984).] Maida and Shapiro's proposal is in the tradition of explicitly bringing in concept-like intensions as a basis for the representation of propositional attitudes. Other proposals in this general line are McCarthy (1979), Creary (1979) and Barnden (1983). My objections to the Maida and Shapiro scheme are probably not fatal; but it is precisely because of the scheme's importance and promise that it is worthwhile to point out types of theoretical inelegance and pragmatic awkwardness that could have been avoided.

Maida and Shapiro place much emphasis on the idea that the nodes in their networks do not follow the usual line of representing "extensions", by which they mean "ordinary" sorts of entities like people, numbers, etc.; rather, nodes "represent" intensions. Intensions can be, for instance, concepts like that of Mike's telephone number as such (where the "as such" indicates that the concept in some sense includes the characterization of the number *as* the telephone number of Mike), or a propositional concept like that of Bill dialling Mike's telephone number. The question now is: which of the two types of representation are they appealing to? If their John node merely represents X's John-concept in the ambassadorial sense, then it *denotes* John. But Maida & Shapiro (1982: p.296) *say explicitly that they are departing from the idea of nodes denoting ordinary objects like John*. So, we have to assume that they mean that the John node actually denotes the John *concept* via the scheme's semantic function, and does not denote John himself — more generally, that the universe of discourse of their scheme is just the world of their chosen agent X's concepts, so that the semantic function of the scheme maps nodes to such concepts.

On the other hand, Maida and Shapiro want their networks to be used by artificial cognitive agents, not just to serve as a theoretical tool, and there are strong signs that Maida and Shapiro really think to an important extent of their network nodes as *ambassadorially* representing intensions. 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 that 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.

Maida and Shapiro's use of the verb "to model" is also suspect. They say in (1982: p.296) that they want their networks to model the belief structures of cognitive agents. It is this desire that is the motivation for their adopting the "intensional" view whereby nodes represent intensions. They failed to notice that a conventional network that ambassadorially represents agent X, but denotes ordinary things in the world, *does* model the belief structure of X. Such a network models the agent X ambassadorially while modelling the agent's world in the denotational sense.

We now go on to see what problems arise in the Maida & Shapiro scheme that are linked to the ambassadorial/denotational confusion and to the distinctions between views (A) and (B) in Section I. The treatment here is necessarily brief. Further details can be found in (Barnden, 1985).

The proposition that John dials Mary's telephone-number could appear in a Maida & Shapiro network in the way shown in Figure 1. The network is associated with a particular cognitive agent X ("the system"). (For convenience, I simplify the form of the networks in harmless ways, and adopt a diagrammatic notation slightly different from the one used by Maida and Shapiro.) The John, Mary, dial and tel-num nodes in Figure 1 denote X's *concept of* John, X's *concept of* Mary, X's *concept of* dialling, and X's *concept of* the telephone-number relation. The "head" node D denotes the proposition, which is itself a concept or intension, and has a truth value as its extension. The node MTN denotes the concept of Mike's telephone number as such.

A. Non-Uniform Dereferencing

Our first difficulty is to do with the semantics of the network fragment in Figure 2, which shows the way the proposition that Bill believes that John is taller than Mary could appear. Rather than trying to appeal to a precise semantics, I shall use an informal, simplified approach. The proposition, denoted by the node B in Figure 2, 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

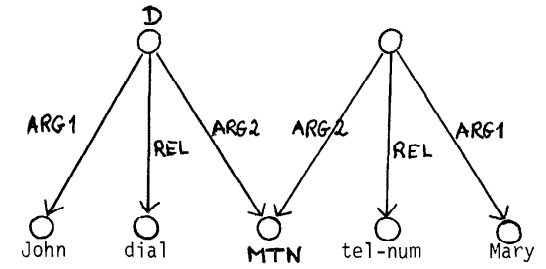

Figure 1

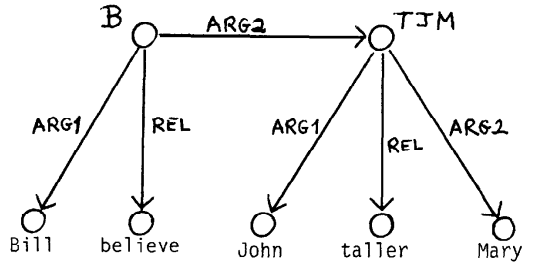

Figure 2

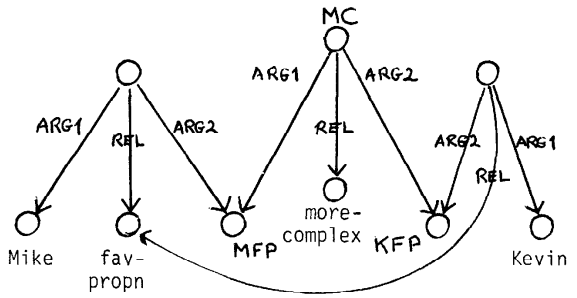

Figure 3

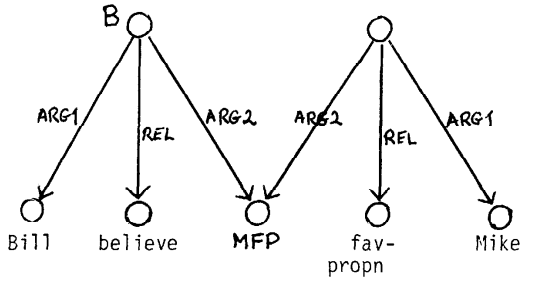

Figure 4

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 theoretical and practical drawback. Theoretical, simply because it is a complication, and it has no analogue in some other powerful concept-based propositional attitude representation schemes, such as that of Creary (1979). Practical, because it forces processing mechanisms that act on the networks to 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 network fragments into natural-language statements. In the Figure 2 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).

Remember that, despite their view of their networks as denoting conceptual structures in agent X, Maida and Shapiro also seem to propose that an AI program could *use* their networks in dealing with the world. Such a program, then, might need the language-generating mechanism of the previous paragraph.

B. Descriptions of Propositions

The complication of having certain argument positions of certain relationships be of type "do-not-dereference" is unwelcome, but perhaps tolerable. However, consider now 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), nor indeed in most studies of propositional attitudes, whether in Philosophy or AI. I am *merely* using the "favourite proposition" function as a way of generating simple examples. Other sorts of mundane descriptions of propositions would be more important in practice (e.g. the description "the belief he expressed yesterday").

Figure 3 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 as such*. This is by analogy with node MTN of Figure 1. Thus, MFP denotes a concept whose extension is a (propositional) concept. Then, in saying what the proposition denoted by node MC in Figure 3 is about, we must dereference MFP and KFP. The relation "more-complex-than" is like the relations "taller-than" and "to dial" in that it does not have do-not-deref argument-positions.

The Figure 3 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 Figure 4. In saying what the proposition denoted by node B states, we *do* now have to dereference the concept denoted by the ARG2 node (i.e. MFP) in the belief structure, in contrast to the case of Figure 2. For, it isn't that Bill believes *the concept of* Mike's favourite proposition; rather, he believes that proposition itself. 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 Figure 2 difficulty, is to refrain from dereferencing just those concepts denoted by "proposition head nodes" — nodes that send out REL arcs. This rule specifies that the MFP and KFP concepts in Figs. 3 and 4 should be dereferenced, but not the TJM concept in Figure 2. But we now get into trouble with the proposition that Kevin's favourite proposition is that John is taller than Mary. The structure in Figure 5 is not satisfactory. The problem is that TJM denotes the proposition that John is taller than Mary, but equally, by analogy with the Figure 3 example, it should denote a concept

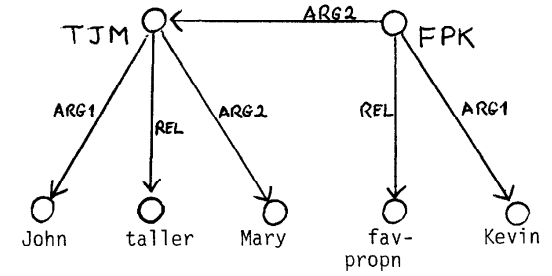

Figure 5

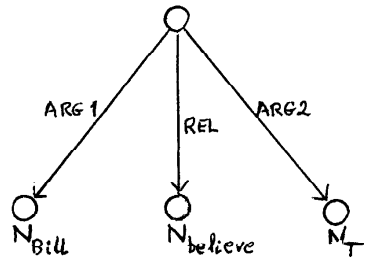

Figure 6

of some proposition, because it is the ARG2 node of a favourite-proposition-of predication.

In Barnden (1985) I comment on the difficulties that would arise in an attempt to avoid the above problems by deploying an explicit dereferencing operator, or by having FPK's ARG2 node (see Figure 5) go to a node different from TJM but related to it by the Maida and Shapiro EQUIV facility.

C. The Source of the Difficulties

Although Maida and Shapiro tried to design a scheme in which the denoting expressions always denote concepts, they were strongly though unintentionally influenced by the more typical idea of a scheme that ambassadorially represents concepts. They correctly realized that a representation scheme able to cope with the representation of agents' propositional attitudes can be achieved by having representational items that denote intensions (concepts with the form, for instance, of propositions and definite descriptions). The same basic idea has appeared in other systems, such as the predicate-logic schemes of McCarthy (1979) and Creary (1979). However, each of these other schemes was *explicitly* proposed for *use by* a cognitive agent, with the denoting items in the scheme denoting people, numbers and so on as well as intensions.

Assume in fact that we were to design a semantic-network scheme along these more conventional lines. Suppose that the agent X using the scheme is entertaining the proposition that Bill believes that John is taller than Mary. Agent X could then contain a cognitive structure P whose semantic network formulation or abstraction is depicted in Figure 6. Notice carefully that the node labelled "N_{Bill}" denotes Bill himself now, not an X-concept of him, so that it is unlike the Bill nodes in previous figures. Equally, node $N_{believe}$ denotes the relation of believing itself, not an X-concept of it. On the other hand, the node N_T is akin to a Maida and Shapiro proposition node in that it denotes the proposition T that John is taller than Mary. I leave other details vague, in particular the means by which node N_T denotes T and the semantics of the topmost node in the figure. Let us assume, then, that nodes N_{Bill} and N_T are the *abstractions from* some cognitive structures C_{Bill} and C_T in agent X, where we say that C_{Bill} and C_T are concepts that X has of the person Bill and the proposition T.

Now suppose that we choose to design a network scheme for *describing* agents such as X. Then we need a node N'_{Bill} to denote X's concept C_{Bill}. This node is therefore similar to the person nodes that we have seen in examples of the Maida and Shapiro system. But, similarly, we also need a node N'_T to denote X's concept C_T of the proposition T. This node is *not* like the Maida and Shapiro node TJM in Figure 5. Node TJM denotes the proposition T itself, whereas our N'_T denotes C_T.

If Maida and Shapiro wanted to build a scheme whose nodes denoted only intensions, they should perhaps have built a scheme containing nodes like N'_T. That they did not can be explained that what they really had at the back of their minds in considering propositions like "Bill believes that John is taller than Mary" was *an ambassadorial representation structure such as the one in Figure 6.* The trouble is that in adopting their intensional view they "lifted Bill up by one intensional level" [they introduced a node denoting an intension for Bill] but failed to lift proposition T up by an intensional level [introduce a node denoting an intension for that proposition].

It is the systematic confusion of U with C_U for propositions U that is at the root of the difficulties for Maida and Shapiro. We can see this by supposing that in Figure 2 node TJM is now taken to denote a *concept of* the proposition T that John is taller than Mary, rather than T itself. Then in determining what the structure in the figure is saying we have to dereference both argument nodes of node B, not just one. This dereferencing in belief subnets then removes the difficulties we encountered with the example of Figure 4.

III SOME SUBTLE IMPUTATIONS

Creary (1979) has proposed a neo-Fregean way of using logic to represent propositional attitudes. The proposal is a development of one by McCarthy (1979). In the intended interpretation of Creary's system, terms can denote propositional and descriptional concepts as well as "ordinary" things.

Assume that agent X is using Creary's scheme as a represention medium and translates inputed English sentences into expressions in the scheme. Then X will *systematically impute probably-incorrect conceptual structures to other agents,* in a subtle way. Suppose X receives the sentence

((S1)) Mike believes that Jim's wife is clever.

The simplest Creary rendering of this sentence (reading it in a "de-dicto" fashion) is

((C1)) `believe(mike, Clever(Wife(Jim))).`

The symbols `jim` and `mike` denote the people Jim and Mike. The symbol `Jim` denotes a particular (apparently standard) concept of the person Jim. The symbol `Wife` denotes a function that when applied to some concept c of some entity delivers the (descriptional) concept of "the wife of [that entity as characterized by c]" *as such.* Thus the term `Wife(Jim)` denotes a complex, descriptional concept. The symbol `Clever` denotes a function that when applied to some concept c of some entity delivers the propositional concept of "[that entity, as characterized by c] being clever", as such. The `believe` predicate is applied to a person and to a propositional concept.

Note that the intuition underlying the use of the `Clever` and `Wife` functions in the formula is that Mike's belief is in some sense couched in a direct way in terms of wife-ness and cleverness. Notice on the other hand that the formula embodies

no claim that *Mike* conjures with entities representing the functions denoted by the Creary symbols **Clever** and **Wife**. The Clever and Wife functions are (so far at least) merely tools *the agent X* uses to "mentally discuss" Mike.

However, suppose now that X inputs the sentence

((S2)) George believes that Mike believes that Jim's wife is clever.

A major Creary interpretation of (S2) is

((C2)) `believe(g, Bel(Mike, Clever$(Wife$(Jim$))))`

Intuitively, this says that George has in his belief space something that says what (C1) says. Our point hinges on the nature of the concepts denoted by the second-order concept terms in the new formula.

The symbol `Jim$` denotes a (standard) concept of a concept of Jim. The symbol `Bel` denotes the function that, when applied to a concept of a person and a concept of a propositional concept, delivers the propositional concept of that person (so characterized) believing that proposition (so characterized). The term `Wife$(Jim$)` denotes a concept of the concept denoted by `Wife(Jim)`. In fact, the concept denoted by `Wife$(Jim$)` *explicitly involves the idea behind the concept-construction function that is denoted by* `Wife`, just as the concept denoted by `Wife(Jim)` explicitly involves the (wifeness) idea behind the wife-delivering function denoted by **wife**. Similarly, the term `Clever$(Wife$(Jim$))` denotes a concept C of the propositional concept denoted by `Clever(Wife(Jim))`, where C *explicitly involves the the idea behind the concept-construction function that is denoted by* `Clever`.

The trouble then arising is that we must conclude that, intuitively, (C2) conveys that *George's belief is couched in terms of the* `Clever` *and* `Wife` *concept-construction functions*. This is analogous to (C1) conveying that Mike's belief is couched in terms of the wife-of function and cleverness property. We now see that (C2) is a ***deviant*** interpretation of (S2), whereas Creary makes it out to be a ***major*** plausible possibility. (C2) is deviant because it takes George to be conjuring with concept-construction functions that no-one except a theoretician (e.g. Creary) could normally be expected to conjure with.

It is convenient to sum up this phenomenon as being an *imputation*, to cognitive agents, of features of X's particular method of describing cognitive agents. X's method uses concept-construction functions like Clever and Wife; and X, probably incorrectly, imputes the use of these functions to other cognitive agents. This observation about Creary's concept-construction functions was implicit in a primitive form in (Barnden, 1983).

The imputation problem in Creary's system could have practical importance, albeit at a sophisticated level, as I show in (Barnden, 1986). I demonstrate there that analogous imputation issues crop up in many types of scheme — for instance, quotational schemes somewhat on the lines of (Burdick 1982, Perlis 1985, Quine 1981), and the B_r-based situation-semantics scheme of Barwise & Perry (1983), if this is naturally extended to deal with nested attitudes.

The above comments have had to be very brief. The matter is portrayed in much greater detail and in a wider context in (Barnden 1986). I hope the comments made are enough to suggest that the importance of the imputations varies greatly according to whether or not the representation scheme at hand is meant to be used as a theoretical tool for the description of propositional attitudes. If this is not the intention — so that the scheme is instead *just* meant to be used by a cognitive agent as a base for its normal cognitive processing — then the imputations may be tolerable, since the inaccuracy of representation that they embody may not prevent the agent reacting appropriately to its environment in most cases. That is, the representation scheme could still be adequate to a heuristically acceptable degree. If, however, the representation scheme is meant to be used as a theoretical tool, then the theorist should be very aware that the scheme is introducing imputations, or, if you like, is slipping in important theoretical assumptions about the *psychology* of agents by a back door. These assumptions may be over and above the ones the theorist was aware of making.

IV CONCLUSION

To tie everything together: we saw that the importance of imputations inherent in a scheme depends greatly on whether the scheme is merely meant to be *used* by agents or is meant to be a *theoretical tool* for describing agents; and the critique of the Maida and Shapiro system shows that an intended view of a scheme as *describing* an agent (perhaps because the scheme is being used as a theoretical tool) can be unwittingly and deleteriously affected by a view of the scheme as being *used* by the agent. The considerations of this paper have led me to devise a propositional-attitude representation scheme that is relatively free of imputation difficulties. A preliminary sketch is to be found in (Barnden, in press).

REFERENCES

Barnden, J.A. "Intensions As Such: An Outline." In *Procs. 8th. Int. Joint Conf. on Artificial Intelligence.* Karlsruhe, W. Germany, August 1983.

Barnden, J. A. "Representations of Intensions, Representations as Intensions, and Propositional Attitudes." Tech. Rep. 172, Computer Science Dept., Indiana University, Bloomington, Indiana, June 1985.

Barnden, J.A. "Imputations and Explications: Representational Problems in Treatments of Propositional Attitudes." To appear in *Cognitive Science.* Tech. Rep. 187(revised), Computer Science Dept., Indiana University, Bloomington, Indiana, April 1986.

Barnden, J.A. "Interpreting Propositional Attitude Reports: Towards Greater Freedom and Control." To appear in *Procs. 7th European Conf. on Art. Int.*, July 1986. In press.

Barwise, J., & Perry, J. *Situations and Attitudes.* Cambridge, Mass.: MIT Press, 1983.

Brachman, R.J. On the epistemological status of semantic networks. In N.V. Findler (ed.), *Associative Networks.* New York: Academic Press, 1979.

Brachman, R.J. & J.G. Schmolze. "An Overview of the KL-ONE Knowledge Representation System." *Cognitive Science* 9:2 (1985) 171–216.

Burdick, H. "A Logical Form for the Propositional Attitudes." *Syntheses* 52 (1982) 185–230.

Creary, L.G. "Propositional Attitudes: Fregean Representation and Simulative Reasoning." In *Procs. 6th. Int. Joint Conf. on Artificial Intelligence*, Tokyo, Japan, August 1979.

Johnson-Laird, P.N. *Mental Models.* Cambridge, Mass.: Harvard Univ. Press, 1983.

McCarthy, J. "First Order Theories of Individual Concepts and Propositions." In J. E. Hayes, D. Michie & L. I. Mikulich (Eds.), *Machine Intelligence 9.* Chichester, England: Ellis Horwood, 1979.

Maida, A.S. & S.C. Shapiro. "Intensional Concepts in Propositional Semantic Networks." *Cognitive Science* 6:4 (1982) 291–330.

Perlis, D. "Languages with Self-Reference I: Foundations." *Artificial Intelligence* 25 (1985) 301–322.

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

Rapaport, W. J., & S.C. Shapiro. "Quasi-Indexical Reference in Propositional Semantic Networks." In *Procs. 10th Int. Conf. on Computational Linguistics*, Stanford Univ., USA, 1984.