175

J. EXPT. THEOR. ARTIF. INTELL. 5(1993)225–235.

Belief spaces as sets of propositions

STUART C. SHAPIRO

Department of Computer Science and Center for Cognitive Science, State University of New York at Buffalo, 226 Bell Hall, Buffalo, NY 14260-2000, USA
email: shapiro@cs.buffalo.edu

Abstract. It is common in the knowledge representation literature for a belief space to be considered to be a set of sentences. Some implications of this stance are examined, and an alternative view, that belief spaces are sets of propositions is developed, and found to be an improvement. This latter view requires that propositions be accepted as entities in the domain of discourse of languages of thought, which, it is argued, accords with commonsense usage. In exchange, the semantics of nested belief expressions is simplified, and certain problems caused by the sentential view are avoided.

Keywords: knowledge representation, beliefs, propositions, language of thought, semantics

1. Introduction

What type of entity is the object of belief? Equivalently, if the 'belief space' of a cognitive agent is the set of entities that the agent believes, what type of entities are the elements of that set? Two answers that have been suggested in the literature of artificial intelligence (AI) are sentences (Moore and Hendrix 1982, Haas 1986, Konolige 1986, Genesereth and Nilsson 1987, Perlis 1988, des Rivières and Levesque 1988, Davis, 1990) and propositions (McCarthy 1979, Charniak and McDermott 1985), with sentences apparently the more popular of the two. In this paper, it will be argued that a belief space is a set of propositions and the implications of this for knowledge representation (KR) formalisms will be examined.

An extensive survey of the philosophical literature on the subject will not be presented, merely a quote from a 1950 paper of Church's (Church 1971):

> For statements such as (1) *Seneca said that man is a rational animal* and (A) *Columbus believed the world to be round*, the most obvious analysis makes them statements about certain abstract entities which we shall call 'propositions' (though this is not the same as Carnap's use of the term), namely the proposition that man is a rational animal and the proposition that the world is round; and these propositions are taken as having been respectively the object of an assertion by Seneca and the object of a belief by Columbus. We shall not discuss this obvious analysis here except to admit that it threatens difficulties and complications of its own, which appear as soon as the attempt is made to formulate

systematically the syntax of a language in which statements like (1) and (A) are possible. But our purpose is to point out what we believe may be an insuperable objection against alternative analyses that undertake to do away with propositions in favour of such more concrete things as sentences (Church 1971, p. 168).

2. Languages of thought and the sentential model

It is taken as uncontroversial that the contents of a belief space are expressed in some 'language of thought', which, in the case of a computerized agent, is a knowledge representation language. When two agents communicate with each other, they do not do so in a language of thought, but in a public communication language (PCL), such as English. The speaker in a dialogue must translate from its language of thought into the public communciation language in order to say something to the hearer, who, in turn, must translate from the public communication language into its language of thought in order to understand what the speaker is saying. These two translation steps are familiar to everyone who has worked on the natural language understanding and generation problems.

According to the approach that a belief space is a set of sentences, the belief space of a computerized agent is a set of sentences of its KR language, and the agent is taken as believing the sentences in that set. Notice that, in this way of speaking, belief is a relation between an agent and a sentence. If, in some KR language, L1, *John*1 is an individual constant denoting some person named 'John', and *Tall* is a predicate constant denoting the set of tall people, then the L1 sentence *Tall*(*John*1) would be true just in case the person denoted by *John*1 is in the set of tall people, i.e. if he is tall. If, moreover, *Believes* is a binary relational constant in our metalanguage (say ML1) denoting the relation that exists between an agent and a sentence when the agent believes the sentence, *Oscar* is an individual constant of ML1 denoting some computerized agent that expresses its beliefs in L1, ' is a quotation symbol in ML1 that maps a sentence of any language into an ML1 expression denoting that sentence, and *Belspace* is a function symbol of ML1 that denotes a mapping from an agent into its belief space, then *Believes*(*Oscar*, '*Tall*(*John*1)) is a sentence of ML1 that asserts that the sentence *Tall*(*John*1) is a member of *Bellspace*(*Oscar*)—that is, that *Oscar* believes that John is tall.

3. Propositions as entities

Note that there are expressions in any language that denote entities of the domain of discourse. These expressions constitute the set of terms of the language, as opposed to the set of sentences of the language, which are usually taken as denoting truth values. The term *John*1 of L1 denotes some person just as the term *Oscar* denotes, in ML1, some agent. Similarly, the ML1 term '*Tall*(*John*1) denotes a sentence of L1. Thus, the domain of discourse of ML1 contains sentences as entities, as well as agents.

The commonsense world that constitutes the domain of discourse of natural languages includes propositions as well as people, sentences, and lots of other things. Some people call propositions 'beliefs' when they are stressing propositions that are believed, and call them 'truths' when stressing propositions that (they believe) are true. For example, in the US Declaration of Independence, we find,

> We hold these Truths to be self-evident, that all Men are created equal, that they are endowed by their Creator with certain unalienable Rights, that among these are Life, Liberty, and the Pursuit of Happiness

and in Lincoln's Gettysburg Address we find,

> a new nation, conceived in liberty and dedicated to the proposition that all men are created equal,

In both documents, we find references to the proposition that all men are created equal. Like other entities in the domain of discourse, propositions can have properties (the proposition that all men are created equal has the property of being the first self-evident truth mentioned in the Declaration of Independence) and can have multiple, though co-extensional, intensions (*John did not know that the proposition that Lincoln said the nation was dedicated to was the first listed as a self-evident truth in the Declaration of Independence.*).

4. A propositional model of belief spaces

Notice that when discussing propositions, it is natural to use a notion of belief that is a relation between an agent and a proposition (*Sam believes the proposition that the founders didn't really believe the proposition that all men are created equal.*). Thus, the author's alternative to the notion that belief spaces are sets of sentences is that belief spaces are sets of propositions. An agent can express its beliefs as terms in its language of thought. Let's redo our examples of L1 and ML1 above as examples of L2 and ML2: *John*1 is an individual constant of L2 that denotes some particular person named 'John'; *Tall* is a function symbol of L2 that denotes the mapping from a person to the proposition that that person is tall; *Tall*(*John*1) is, therefore, a term of L2 that denotes the proposition that the person denoted by *John*1 is tall; *Oscar* is an individual constant of ML2 that denotes some agent that uses L2 as its language of thought, *Believes* is a function symbol of ML2 that denotes a mapping from an agent and a proposition into the proposition that that agent believes that proposition; *Belspace* is a function symbol in ML2 that denotes a mapping from an agent into the set of propositions that that agent believes. Temporarily, let ML2 include all of L2 so that an expression of L2 denotes in ML2 exactly what it does in L2; *Believes*(*Oscar*, *Tall*(*John*1)) is then a perfectly good term of ML2 denoting the proposition that *Oscar* believes the proposition that *John*1 is tall, that is, that *Tall*(*John*1) is a member of the set *Belspace*(*Oscar*).

Now, it is certainly possible that two expressions, each in a different language, denote the same entity in the domain of discourse that is common to both languages. Therefore, instead of incorporating L2 into ML2, let us say that *IsTall* is a function symbol of ML2 that denotes exactly the same mapping that *Tall* denotes in L2, and that *Jack*2 is an individual constant in ML2 that denotes the same person that *John*1 denotes in L2. Therefore, *IsTall*(*Jack*2) denotes in ML2 exactly the same proposition that *Tall*(*John*1) denotes in L2, and *Believes*(*Oscar*, *IsTall*(*Jack*2)) is a pure ML2 term denoting the proposition that *Oscar* believes the proposition that *Jack*2 is tall. That is, *Believes*(*Oscar*,*IsTall*(*Jack*2)) asserts that the proposition that is the value of $IsTall(Jack2)_{ML2}$ and of $Tall(John1)_{L2}$ is a member of $Belspace(Oscar)_{ML2}$.

The first benefit we see to having a belief space be a set of propositions rather than a set of sentences is the elimination of the need of quotation for expressing the proposition that a particular agent has a particular belief.

Since ML2, as now sketched, is no longer a *meta*-language, let us rename it L3 for future reference.

5. Knowledge base vs. belief space

At any time, a cognitive agent has conceived of more propositions than (s)he believes. Some are the objects of other propositional attitudes, some are objects of nested belief propositions, etc. We need to distinguish the set of KR expressions the agent has constructed from the set of propositions the agent believes. Since we are using *belief space* to refer to the latter, let us use *knowledge base* (*KB*) to refer to the former. Some of the expressions in an agent's KB denote propositions in his/her belief space. These are notated in some way, but we would not want that technique to entail that, whenever an agent believes *P*, (s)he necessarily believes that (s)he believes *P*. (See, for example, the use of the assertion tag in SNePS 2.0, Shapiro and Rapaport 1992, p. 251.) We will say that a KB expression is *asserted* when it is notated in the appropriate way that the proposition is in the agent's belief space (cf. Shapiro 1991).

6. Nested beliefs

Next, let's consider the case of nested beliefs—allowing one agent to have beliefs about the beliefs of other agents. For this purpose, let's introduce the computerized agent Cassie, and let's first consider having Cassie use ML1 as her KR language, while Oscar uses L1 as his KR language. For Cassie to believe that Oscar believes that John is tall, Cassie must have in her belief space the sentence *Believes*(*Oscar*, *'Tall*(*John*1)). The problem, however, is that *Tall*(*John*1) is a sentence of L1, and Cassie does not know L1—she only knows ML1 and some PCL. Therefore the sentence *Believes*(*Oscar*, *'Tall*(*John*1)) does not mean for Cassie that Oscar believes that John is tall, but that Oscar believes some gibberish. Cassie is incapable of relating the beliefs that she believes Oscar to hold to any beliefs of her own. Let's examine the alternatives:

Cassie expresses Oscar's beliefs in the PCL. But if we retain the semantics of $Believes_{\mathrm{ML1}}$, this would mean that Oscar stores in his belief space sentences of the PCL, contradicting the entire point of KR languages. Alternatively, we could change the semantics of $Believes_{\mathrm{ML1}}$ so that its second argument is a sentence of the PCL, and its meaning is that the agent believes the translation of the sentence into its language of thought. But this would require a translation of a PCL sentence into ML1 every time Cassie thinks about what Oscar believes. It is, however, essentially the proposal of Moore and Hendrix (1982), who only discuss the semantics of PCL sentences of the form 'A believes S', not language of thought sentences of the corresponding form.

Cassie expresses Oscar's beliefs in her own language of thought. But this would be a claim that Oscar has in his belief space a sentence of someone else's language of thought—again contradicting our assumptions. (However, this is close to what will be proposed as a final solution.)

Cassie knows Oscar's language of thought. This would entail that every agent knows every other agent's language of thought. This is patently wrong, but could be patched by:

Cassie uses the same language of thought as Oscar. This (or the next variant) seems to be the implicit assumption of people who advocate sets of sentences as belief spaces, and entails that there is only one language of thought used by all agents. It seems that this is also patently ridiculous. Not only would it mean that every agent uses the same symbol for every common noun and predicate, it also means that every agent uses the same individual constant for every individual. If we use symbols of the sequence *John*1, *John*2,. . . for the people named 'John' that we meet, we would all have to meet the same people in the same order, or somehow magically reserve some *John* symbols for people we don't know. A variant that has been proposed by some researchers is:

All agents use the same language of thought except for the denotations of individual constants, function symbols, and predicate symbols. This theory, explicitly proposed by Moore and Hendrix (1982), seems reasonable, at least for agents of the same species, because it seems reasonable to suppose that the syntax of the language and the meaning of its interpreted symbols (e.g. logical constants) are innate. But it seems to pose great problems for cross-species communication and for human–computer communication (which some might see as a virtue), and we're still left with the problem of formulating in one sublanguage an expression denoting a sentence in another sublanguage. One proposed aid to this translation is the use of an explicit *denotes* predicate, where $denotes(a,s1,s2)$ would mean that agent a denotes by $s1$ the same entity that I (the agent in whose language of thought this sentence is expressed) denote by $s2$ (cf. Maida 1991). However, this seems to require that Cassie can somehow examine Oscar's mind and see how he uses his symbols, which also seems patently ridiculous. That ridiculous requirement might be avoided with the use of existential quantifiers (Maida 1991, p. 334). To see how this would go, let's add to ML1: the predicate symbol *IsTall* to denote the set of tall people; the individual constant *Jack*1 to denote the same individual denoted by *John*1 in L1; the function $apply('P,'x)$, which takes a predicate symbol $'P$ and an expression $'x$, and denotes the sentence $'P(x)$ (cf. Davis 1990, p. 369). So now, for Cassie to believe that Oscar believes that John is tall, Cassie would have to have in her belief space the sentence

$$\begin{aligned}&\exists P \exists x(denotes(Oscar,P,'IsTall)\\&\quad \wedge denotes(Oscar,x,'Jack1)\\&\quad \wedge Believes(Oscar,apply(P,x)))\end{aligned}$$

This is both horribly complicated, and it hides the propositional content of Oscar's belief. In any case, it will be shown that, by using sets of propositions instead of sets of sentences, this assumption that all agents use the same language of thought is unnecessary.

Two additional arguments against quoted L1 sentences as the second argument of $Believes_{\mathrm{ML1}}$ were given by Jim Davis (personal communication):

(1) How did Cassie come to have this belief about Oscar's beliefs? Either she inferred it from Oscar's behaviour, in which case she would have composed it in ML1, or Oscar or someone else told her, in which case she would have heard it in PCL. In neither case would Cassie have had access to an L1 sentence.

(2) 'Cassie can have beliefs about the "beliefs" of objects which don't even have internal KRs. People routinely reason about the intentions, desires,

and beliefs of inanimate machines, pets, deceased and fictional persons, and supernatural creatures. Surely none of us can claim to have quoted sentences from, say, Clark Kent's KR, in our heads'.

7. Nested belief propositions

Now let's consider the proposal that belief spaces are sets of propositions, that Oscar uses L2 as his KR language, and that Cassie uses L3 as her KR language. For Cassie to believe that Oscar believes that John is tall, she must store the L3 term *Believes*(*Oscar*,*IsTall*(*Jack*2)) in her KB in such a way as to indicate that the proposition it denotes is in her belief space. It seems that all the problems cited above are now solved. Cassie understands the belief she believes Oscar to hold because it is expressed in her own KR language. Her belief is not about the L2 expression Oscar presumably has in his KB, but about the proposition she believes to be in Oscar's belief space as a result of his proper storage of that L2 expression. L2 may be a radically different language than L3, without making a difference. Every agent may use its own, private, language of thought, and use that language of thought to express propositions it thinks other agents believe. The agents will still communicate with each other in the PCL.

It might be argued that *Jack*2 or *IsTall* might mean in L3 something slightly different from what *John*1 or *Tall* mean in L2, because of slight differences in Cassie's and Oscar's beliefs, and that therefore, even if Oscar has *Tall*(*John*1) in his belief space, Cassie is wrong if she has *Believes*(*Oscar*,*IsTall*(*Jack*2)) in hers. This possibility is accepted, but it is taken as a positive feature of this theory, rather than as a negative. It models the disagreements people can have with each other over the subtle aspects of the meanings of words they use in common.

The second benefit we see to having a belief space be a set of propositions rather than a set of sentences is the solution of the trilemma of either having all agents use the same language of thought, or having agents not understand what beliefs they ascribe to other agents, or having PCL sentences used as part of KR belief sentences.

There are two proposals that admit that belief spaces are sets of propositions, yet preserve the syntax of quoted sentences:

(1) Change the semantics of $Believes_{\mathrm{ML1}}$ so that *Believes*(*A*,*P*) means that agent *A* believes the proposition expressed by the sentence *P*. Although this allows propositions into belief spaces, and allows *P* to be expressed in the language of the outer believer (Cassie can represent her belief that Oscar believes that John is tall by *Believes*(*Oscar*,'*IsTall*(*Jack*2))), it does not admit propositions into the ontology in the sense that no term has a proposition as its denotation and no variables range over propositions.

(2) Change the semantics of the quotation symbol ', so that, if *P* is a sentence, '*P* is the proposition expressed by *P*. The main problem with this is that '*x* is syntactically ill formed, so we cannot represent such sentences as 'John believes everything Bill believes' unless we add additional proposition-forming machinery.

Neither of these solutions has the elegance of the solution argued for here, namely that the sentence of the sentential model instead be considered a term that denotes a proposition. The elegance comes from remaining a first-order theory without quotation operators or modal operators.

8. Inclusion of logical constants and logical inference

The ability of the sentential model to represent disjunctions, implications, and other non-atomic sentences in belief spaces, and to use the mechanisms of logical inference on these sentences to generate new beliefs need not be lost by the change to the propositional model. The propositional connectives now become functions from propositions to propositions instead of functions from truth values to truth values, and rules of inference can be formulated and implemented of the form: Whenever an expression of the form $P \Rightarrow Q$ is asserted in the KB, and the subexpression P is also asserted, assert the subexpression Q.

A more formal presentation of the language being proposed here is:

Individual constants: There is a set of individual constants, $a_1, a_2, \ldots$, each denoting an entity in the domain. Among these are a set of propositional constants, denoting propositions. In this paper, we will let $[\![a_1]\!]$ express the denotation of the constant a_i.

Variables: There is a set of variables, $x_1, x_2, \ldots$, that range over entities in the domain. Among these are a set of propositional variables that range over propositions.

Function symbols: There is a set of function symbols, $f^1_1, f^1_2, \ldots, f^2_1, f^2_2, \ldots$. Each function symbol f^i_j denotes a mapping from the Cartesian product of some particular i sets of entities to some particular set of entities. In this paper, we will let $[\![f^i_j]\!]$express the denotation of the function symbol f^i_j. Among the function symbols are a set of propositional function symbols whose range is the set of propositions. Among the proposition functions are these:

- $\neg$: ⟨propositions⟩ $\rightarrow$ ⟨propositions⟩, which maps a proposition ϕ into the proposition that it is not the case that ϕ.
- $\wedge$: ⟨propositions⟩ $\times$ ⟨propositions⟩ $\rightarrow$ ⟨propositions⟩, which maps the propositions ϕ and ψ into the proposition that it is the case that both ϕ and ψ.
- $\vee$: ⟨propositions⟩ $\times$ ⟨propositions⟩ $\rightarrow$ ⟨propositions⟩, which maps the propositions ϕ and ψ into the proposition that either it is the case that ϕ or it is the case that ψ.
- $\Rightarrow$: ⟨propositions⟩ $\times$ ⟨propositions⟩ $\rightarrow$ ⟨propositions⟩, which maps the propositions ϕ and ψ into the proposition that if it is the case that ϕ then it is the case that ψ.
- $\Leftrightarrow$: ⟨propositions⟩ $\times$ ⟨propositions⟩ $\rightarrow$ ⟨propositions⟩, which maps the propositions ϕ and ψ into the proposition that it is the case that ϕ if and only if it is the case that ψ.

The previous four function symbols are usually used with an infix notation (for example, $(\phi \Rightarrow \psi)$) instead of a prefix notation $(\Rightarrow (\phi, \psi))$.

Terms:

1. Every individual constant is a term. A propositional constant is a propositional term.
2. Every expression of the form $f_j^n(\phi_1, \ldots, \phi_n)$, where f_j^n is a function symbol and each ϕ_i is a term whose denotation is in the appropriate set, is itself a (functional) term denoting the entity $[\![f_j^n]\!]([\![\phi_1]\!], \ldots, [\![\phi_n]\!])$. A functional term that denotes a proposition is called a propositional term.
3. If x_i is a variable, and ϕ is a propositional term containing one or more instances of the subterm ψ, and none of these instances of ψ is in the scope of a quantifier $\forall x_i$ or $\exists x_i$, and $\phi\{x_i/\psi\}$ is the result of replacing all those instances of ψ by x_i, then $\forall x_i(\phi\{x_i/\psi\})$ and $\exists x_i(\phi\{x_i/\psi\})$ are propositional terms, $\phi\{x_i/\psi\}$ and all its subterms are said to be in the scope of the quantifier $\forall x_i$ or $\exists x_i$, respectively, and all the occurrences of x_i that replaced occurrences of ψ in $\phi\{x_i/\psi\}$ are said to be free in $\phi\{x_i/\psi\}$. The denotation of a term $\forall x(\phi)$ is the proposition that for every term τ of the appropriate type (that is, $[\![\tau]\!]$ is in the appropriate set for every functional term in which τ would occur as an argument), it is the case that $\phi\{\tau/x\}$, where $\phi\{\tau/x\}$ is a term that results from replacing every occurrence of x that is free in ϕ by τ. Similarly, the denotation of a term $\exists x(\phi)$ is the proposition that for some term τ of the appropriate type, it is the case that $\phi\{\tau/x\}$.

Rules of inference: The rules of inference will be displayed as

$$\begin{array}{c} P \\ \underline{Q} \\ R \end{array}$$

and will mean that if an agent believes $[\![P]\!]$ and $[\![Q]\!]$, then that agent may conclude $[\![R]\!]$. It does not mean that if an agent believes $[\![P]\!]$ and $[\![Q]\!]$, it also does believe $[\![R]\!]$. The latter would cause the agents to be logically omniscient. Given this understanding, the appearance of the rules of inference are the same as for standard FOPC, for example,

$$\begin{array}{c} \phi \\ \phi \Rightarrow \psi \\ \hline \psi \end{array}$$

and

$$\begin{array}{c} \forall x(\phi) \\ \hline \phi\{\tau/x\}, \text{ where } \tau \text{ is any term of the appropriate type} \end{array}$$

9. Equality and propositions

We must consider how propositions interact with equality. After all, we do not want our theory to reintroduce the incorrect inference

$$\begin{array}{c} \mathit{MorningStar} = \mathit{EveningStar} \\ \underline{\neg \mathit{Believes}(\mathit{John}, \mathit{MorningStar} = \mathit{EveningStar})} \\ \neg \mathit{Believes}(\mathit{John}, \mathit{MorningStar} = \mathit{MorningStar}) \end{array} \tag{1}$$

(Note that (1), (2), and (3) are here presented informally for motivation. They will be formalized below.) However, we do want our theory to allow (Barnden 1986)

$$\frac{\begin{array}{c}\textit{That John is taller than Mary is Kevin's favourite proposition.}\\ \textit{Bill believes Kevin's favourite proposition.}\end{array}}{\textit{Bill believes that John is taller than Mary}} \tag{2}$$

and, of course, we want

$$IsTall(Jack2)_{L3} = Tall(John1)_{L2} \tag{3}$$

The principles that make this work are from, or follow from, those in Maida and Shapiro (1982), namely:

(1) Terms of KR languages denote intensional entities.
(2) No two distinct terms of a single KR language denote the same entity (the uniqueness principle), so that there is no overriding principle that one can replace any other in a term and maintain the denotation of the term. However some intensional entities may be coreferential such as the Morning Star and the Evening Star. We'll say $Equiv(e_1,e_2)$ for the proposition that $[\![e_1]\!]$ and $[\![e_2]\!]$ are coreferential. If $Equiv(e_1,e_2)$, e_1 and e_2 may replace each other in a term only when explicitly sanctioned by a rule. *Equiv* is an equivalence relation.
(3) A proposition (an intensional entity) is a function of its constituents (also intensional entities).
(4) The rule

$$\frac{\begin{array}{c}Equiv(P,Q)\\ Believe(a,P)\end{array}}{Believe(a,Q)}$$

is valid.

The '=' of (1) and the central 'is' of the first line of (2) are both coreference, while the '=' of (3) is identity. Therefore, a formalization of (1) is

$$\frac{\begin{array}{c}Equiv(MorningStar,EveningStar)\\ \neg Believes(John,Equiv(MorningStar,EveningStar))\end{array}}{\neg Believes(John,Equiv(MorningStar,MorningStar))}$$

which is invalid; a formalization of (2) is

$$\frac{\begin{array}{c}Equiv(\textit{That John is taller than Mary},\textit{Kevin's favourite proposition})\\ Believes(\textit{Bill},\textit{Kevin's favourite proposition})\end{array}}{Believes(\textit{Bill},\textit{that John is taller than Mary})} \tag{6}$$

which is valid by Principle 4; and $IsTall(Jack2)_{L3}$ is the same proposition as $Tall(John1)_{L2}$ by Principle 3.

10. Truth values as properties of propositions

So far propositions have been characterized only to the extent of saying that they are entities of the domain of discourse, and functions of their participants. It seems necessary, however, to recognize that some propositions are true, and that others are false. To account for this, truth values will be treated as properties of propositions, just as height is a property of a person. Presumably, an agent will believe the proposition *P* whenever (s)he believes that the truth value of *P* is true, but treating truth value as a property of propositions gives us a clean representation of sentences such as, 'John knows whether he is taller than Bill'. For a discussion of this representation, see Maida and Shapiro (1982).

11. Related work

A formulation of proposition-valued terms essentially like the one argued for here was made by McCarthy (1979). Proposition-valued terms are also used in Kowalski (1979), and the suggestion that 'a formula denotes a "proposition"' is made in Charniak and McDermott (1985, p. 322). This author has long used sentence or proposition level formulae as terms in other KR formulae (especially as objects of belief) without quotation operators. For example, see Shapiro (1971, 1979), Maida and Shapiro (1982), Shapiro and Rapaport (1987, 1991, 1992), Shapiro (1991).

Nevertheless, any thought that the use of proposition-valued terms for the representation of nested beliefs is a well-known method is mistaken. For example, both Genesereth and Nilsson (1987) and Davis (1990) discuss the representation of nested beliefs, but only discuss the sentential (including modal) techniques.

Acknowledgements

This work was sponsored in part by DARPA and Rome Laboratory under USAF contract F30602-91-C-0040 via a subcontract from Paramax Systems Corporation.

Comments by William J. Rapaport, members of the UB Center for Cognitive Science, members of the SNePS Research Group, participants of the 1992 AAAI Symposium on Propositional Knowledge Representation, and two anonymous reviewers for KR'92 on earlier drafts of this paper are greatly appreciated.

References

Barnden, J. A. (1986) A viewpoint distinction in the representation of propositional attitudes. In *Proceedings of the Fifth National Conference on Artificial Intelligence*, (San Mateo, CA: Morgan Kaufmann) pp. 411–415.

Charniak, E. and McDermott, D. (1985) *Introduction to Artificial Intelligence* (Reading, MA: Addison-Wesley).

Church, A. (1971) On Carnap's analysis of statements of assertion and belief. In L. Linsky (ed.) *Reference and Modality* (Oxford: Oxford University Press) chapter XI, pp. 168–170. (Reprinted from *Analysis*, **10**(5) (1950), 97–99.)

Davis, E. (1990) *Representations of Commonsense Knowledge* (San Mateo, CA: Morgan Kaufmann).

des Rivières, J. and Levesque, H. J. (1988) The consistency of syntactical treatments of knowledge. *Computational Intelligence*, **4**(1): 31–41.

Genesereth, M. R. and Nilsson, N. J. (1987) *Logical Foundations of Artificial Intelligence* (Los Altos, CA: Morgan Kaufmann).

Haas, A. (1986) A syntactic theory of belief and action. *Artificial Intelligence*, **28**: 245–292.

Konolige, K. (1986) *A Deduction Model of Belief* (London: Pitman).

Kowalski, R. (1979) *Logic for Problem Solving* (New York: North Holland).

Maida, A. S. (1991) Maintaining mental models of agents who have existential misconceptions. *Artificial Intelligence*, **50**(3): 331–383.

Maida, A. S. and Shapiro, S. C. (1982) Intensional concepts in propositional semantic networks. *Cognitive Science*, **6**(4): 291–330. Reprinted in R. J. Brachman and H. J. Levesque (eds) *Readings in Knowledge Representation* (San Mateo, CA: Morgan Kaufmann 1985) pp. 170–189.

McCarthy, J. (1979) First order theories of individual concepts and propositions. In J. E. Hayes, Donald Michie, and L. I. Mikulich (eds) *Machine Intelligence 9* (Chichester, England: Ellis Horwood) pp. 129–147.

Moore, R. C. and Henrix, G. G. (1982) Computational models of belief and the semantics of belief sentences. In S. Peters and E. Saarinen (eds) *Processes, Beliefs, and Questions* (Dordrecht, Holland: D. Reidel) pp. 107–127.

Perlis, D. (1988) Languages with self-reference II: Knowledge, belief, and modality. *Artificial Intelligence*, **34**: 179–212.

Shapiro, S. C. (1971) A net structure for semantic information storage, deduction and retrieval. In *Proceedings of the Second International Joint Conference on Artificial Intelligence* (San Mateo, CA: Morgan Kaufmann) pp. 512–523.

Shapiro, S. C. (1979) The SNePS semantic network processing system. In N. V. Findler (ed.) *Associative Networks: The Representation and Use of Knowledge by Computers* (New York: Academic Press) pp. 179–203.

Shapiro, S. C. (1991) Cables, paths and "subconscious" reasoning in propositional semantic networks. In J. Sowa (ed.) *Principles of Semantic Networks: Explorations in the Representation of Knowledge* (Los Altos, CA: Morgan Kaufmann) pp. 137–156.

Shapiro, S. C. and Rapaport, W. J. (1987) SNePS considered as a fully intensional propositional semantic network. In N. Cercone and G. McCalla (eds) *The Knowledge Frontier* (New York: Springer-Verlag) pp. 263–315.

Shapiro, S. C. and Rapaport, W. J. (1991) Models and minds: Knowledge representation for natural-language competence. In R. Cummins and J. Pollock (eds) *Philosophy and AI: Essays at the Interface* (Cambridge, MA: MIT Press) pp. 215–259.

Shapiro, S. C. and Rapaport, W. J. (1992) The SNePS family. *Computers & Mathematics with Applications*, **23**(2–5): 243–275. Reprinted in F. Lehmann (ed.) *Semantic Networks in Artificial Intelligence* (Oxford: Pergamon 1992) pp. 243–275.