Multiple Agents and the Heuristic Ascription of Belief.

Yorick Wilks and Afzal Ballim

Rio Grande Research Corridor, Computing Research Laboratory, New Mexico
State University, Box 3CRL, Las Cruces, NM 88003, USA.

ABSTRACT

A method for heuristically generating nested beliefs (what some agent believes that another agent-believes ... about a topic) is described. Such nested beliefs (points of view) are esential to many processes such as discourse processing and reasoning about other agents' reasoning processes. Particular interest is paid to the class of beliefs known as *atypical beliefs* and to intensional descriptions. The heuristic methods described are embboddied in a program called *ViewGen* that generates nested viewpoints from a set of beliefs held by the system.

1. Introduction

We describe a computational model of beliefs for natural language understanding, planning and reasoning. In particular we are interested in generating complex belief structures known as points of view (what an agent believes about some topic). The model is based upon prior work reported in Wilks and Bien (1979; 1983).

A Prolog program (*ViewGen)* for generating points of view is discussed, that ascribes belief using default reasoning.

2. Background

A point of view is a set of beliefs held about some topic by an agent. A nested point of view is a set of beliefs about some topic that one agent believes another agent believes ... another agent holds. For example agent A may believe that agent B believes that agent C has some set of beliefs about world hunger. We use a shorthand (as in figure 1) to represent a nested point of view.

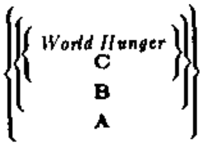

Figure 1. **A's view of B's view of C's view of World Hunger.**

2.1. Relationship To Other Work

The aim of this work has been, since that of Wilks & Bien (1979, 1983), to construct a formalism and programs that capture the heuristic belief ascriptions that individuals actually perform in the process of understanding and participating in dialogue: that is to say, contentful, concrete beliefs and not merely meta-beliefs about the reasoning architecture of others, activities we suspect are rarely, if ever, undertaken in fact. Our concern has been less with the powers of particular notations and proofs of their adequacy (as has been the central feature of the work of Creary 1979, Moore & Hendrix 1979, Konolige 1985, and Attardi & Simi 1984), than with the *content of* belief ascription.

In that sense, our work has been closer in spirit to that of Perrault and his colleagues (e.g. Perrault & Allen 1980), though without their (then) commitment to the language of speech act theory and, most importantly, without their key assumption that the partitions within nested beliefs are all present at the beginning of the belief ascription procedures. Our work makes no such assumption: nested beliefs are not merely accessed but constructed and maintained in real time, a position we find both computationally and psychologically more plausible. The Gedanken Experiment here is to ask yourself if you already know what Mr Reagan believes the Ayatollah believes about Col.Gaddafi. How plausible is it that you have already pre-computed such nested belief spaces?

The work of Maida (1980) clearly shares many of the concerns of the current work: his diagramatic representation of nested beliefs are isomorphic to those of Wilks & Bien (1979) and Shadbolt (1983). He shares, for example, our concern with intensional rather than explicitly extensional representations, as does the SNePS group (Shapiro & Rapaport, 1986). However, his concern is still the problem of shared reasoning strategies between believers and how, for example, you could establish that a dialogue partner also used *modus ponens.* We argue, on the contrary, that this phenomenon is best handled by general default assumptions, as are the concrete contents of belief.

3. Heuristics for Constructing Points of View: a contextual focus of beliefs.

The question arises of how to select the initial beliefs about a topic. Our solution to this problem is essentially simple. The system's belief space is divided into a number of topic specific sub-spaces. These sub-spaces may be thought of as a less permanent version of *frames* (Minsky, 1975) or more suitably in terms of (Wilks, 1977) as *pseudo-texts* (henceforth PTs). In effect, a PT is a set of unsorted, unrefined items of knowledge.

These PTs are general items and are not only stored for individual human beings, but also for groups of humans, objects, and abstract ideas. Their hierarchial and inheritance relations are discussed in Wilks (ibid).

Our method of restricting the initial set of beliefs is to form this set from a very small number of PTs; for example if the system is a medical diagnostician, being used to advise a patient on thalassemia, then the initial set of beliefs may be just those concerning thalassemia. We term the initial belief set the *contextual focus.* We now turn to the process of constructing a point of view from a given contextual focus.

3.1. Generation of Points of View

Points of view are represented by structures known as environments; environments are structures that consist of a sequence of agents and a set of beliefs (e.g. (<a,b>,B) represents that a believes b believes the set of beliefs B).

The process of generating an environment may be regarded as a decision mechanism that ascribes beliefs from one environment to another.

Assume that our contextual focus is as follows:

Figure 2. A Contextual Focus of Belief.

This represents the beliefs held by the sytem that an agent known as the User dislikes another agent known as Frank; that the system likes Frank; that the Earth is round; and that the system believes that Frank believes the Earth is flat. Given this set of beliefs how do we generate what the system believes Frank believes about the contextual focus?

The proposal (Wilks & Bien, 1979;1983) is that a nesting should be generated by a form of default reasoning, using a default rule for ascription of beliefs. The default ascriptional rule is to assume that one's view of another person's view is the same as one's own *except where there is explicit evidence to the contrary.* Applying this rule to (figure 2) we get the system's view of Frank's view (figure 3).

Figure $. The Result of Applying the Default Rule to Figure 2.

Let us examine this example closely. The belief of the system in (figure 2) that the system likes Frank has been ascribed into the PT for Frank, i.e. it is now a belief of the system that Frank believes that the system likes him. The same is true of the system's belief that the user dislikes Frank. However the system's belief that the Earth is round has been superceded by the already existent belief, that the system believes Frank holds, that the Earth is flat. The result is (figure 3) and we may now reason with the environment of the system's views of Frank's views[1], (that we shall call *Sy»tem*$_{Fni)k}$).

Nested points of view are generated using a method called pushing down environments[2], that recursively ascribes beliefs from one environment to another to achieve the nesting.

Suppose that we wish to construct the system's view of the user's view of Frank's view about some context, (we shall write this as *Sy8tem*$_{UMtliv}$^(*'context)* as a shorthand). To construct this we firstly construct the system's view of the user's view and push the system's view of Frank's view down into this. For example presume that we have the contextual focus shown in (figure 4).

Figure 4- A Contextual Focus about the Shape of Earth.

We construct the system's view of the user's view according to our default rule of ascription to get (figure 5).

Figure 5. The System's View of the User's View of the Shape of Earth.

We have available the system's view of Frank's view that is shown in (figure 6).

Figure 6. The System's View of Frank's View of the Shape of Earth.

Now, pushing the environment of (figure 5) down into the environment of (figure 6) results in the environment shown in (figure 7).

Figure 7. The System's View of the User's View of Frank's View of the Shape of the Earth.

To construct a deeper nesting we continually apply this method. So to construct *System*$_{v,er}$ we push the system's view down into the user, push the result of that down into Frank, and finally push the result of that operation down into the system. The mechanism is described in more detail in (Ballim, 1986).

4. Counter Evidence to the Default Ascription of Belief

Thus far we have shown how the default rule can ascribe belief to an agent if there is no further information (as in ascribing the belief that Sally dislikes Frank, to Sally in figures 2,3) and how the rule can be overridden by an *a priori* belief that contradicts the belief that we are attempting to ascribe (as in figures 4,5). We now consider more complex cases.

4.1. Atypical Beliefs

We identify a special class of belief (known as *atypical belief)* that requires a rule opposite to that of default ascription.

An atypical belief is a belief that is held by an agent but would not be held generally by other agents. The class of atypical beliefs covers such areas as self knowledge, secrets, expertise and knowledge of uncommon domains (such as the believer's hobbies, skills, etc.).

So, for example, the belief that the Earth is flat is atypical. An important point must be made here. In terms of a specific agent the foregoing definition is insufficient. For a belief to be considered atypical, with respect to an agent, the agent must believe the belief to be atypical, i.e., 1 may believe you to have an atypical belief, however, I may also believe that you think it is a typically held belief. So while I believe it to be atypical, you believe it to be typical.

For the class of atypical beliefs the rule should be not to ascribe unless one has explicit evidence to justify ascribing the belief. The problem is one of representing and handling a wide range of types of atypical beliefs. The introduction of meta-beliefs is one possible solution, because meta-beliefs enable explicit representation of atypical beliefs (i.e., for atypical belief p, have the meta-belief *atypical(p)).* Due to the wide range of atypical beliefs, and to the problems that they pose, we use a special representation to handle them.

4.2. Intensional Descriptions and Atypical Belief

McCarthy and others have suggested that lambda expressions be used to represent knowledge of values. In (Wilks, 1986) it is proposed that expertise may be expressed within a system by use of lambda expressions with restrictions on the capable evaluators of each such lambda expression. Knowing or having the belief represented by the lambda expression means that the agent is capable of evaluating the expression. So the representation for a cure for tuberculosis is:

(CURE-FOR Tuberculosis) BE
(X(ar) *{CURE-FOR- TB x))<MD*>*

where the only capable evaluators are those known to be MDs (medical doctors). These lambda expressions can be viewed as intensional expressions, i.e. expressions that return the value of some intensional description (cf. Barnden 1983; Maida, 1983). Beliefs involving such intensional expressions are generally atypical as shown by the problems of iterated propositional attitudes that involve intensional descriptions (Creary, 1979; Maida, 1983).

We feel that the problems of expressing beliefs about intensional descriptions are so closely aligned with those of atypical beliefs in general that a representation for differing beliefs about the referent of intensional descriptions can be used for atypical beliefs.

4.2.1. A Taxonomy of Meta-Beliefs about Values

The form of representation suggested in (Wilks, 1986) allows expressions that can only be evaluated by specific agents or classes of agents. This is a first step towards representing atypical belief, however, it is insufficient.

Consider the major factor that makes a belief atypical; some agent believes that a belief held by another agent is not commonly held. A belief about another belief is known as a meta-belief. It is the meta-beliefs about a belief that mark a belief as atypical. An atypical belief can have a large number of meta-beliefs associated with it.

These meta-beliefs can often be classified according to the relation that they define between an agent and the belief with which they are concerned. We propose a taxonomy of meta-beliefs (beliefs about what agents believe is the value of some expression) as shown in (figure 8).

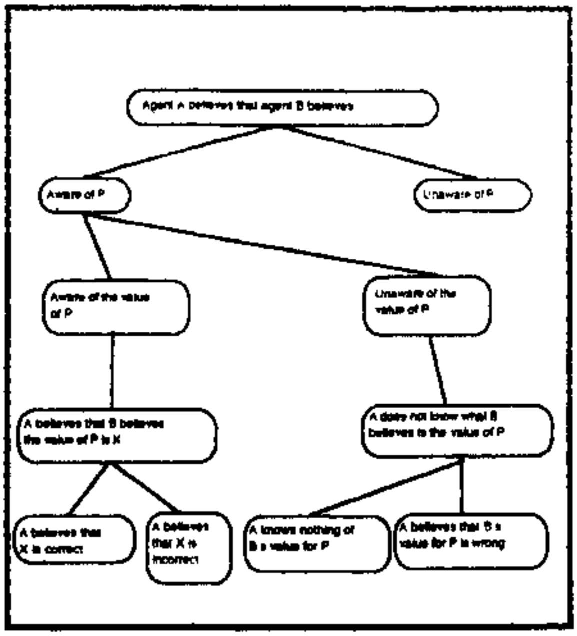

Figure 8. A Taxonomy of Meta-Beliefs about Values

Consider the standard belief about an agent's phone number (call this belief P), i.e., the belief is that the agent's phone number is some value. An agent (A) may have a number of meta-beliefs concerning this belief P. In figure 8 we see that agent A may believe another agent (2?)[3] to be either aware or else unaware of the belief (P). Somebody who does not know of John can hardly have beliefs about John's phone number, and so that person should be unaware of any beliefs about John's phone number.

Further types of meta-belief are possible if A believes B is aware of P. A may believe that B is unaware of the value of P. For example, I may believe that you know that John has a phone number but that you do not know what his phone number is.

If agent A believes that agent B does know the value of P, then A may either know what B believes is the value, or else A may not know what B believes is the value of P. I may know that you believe John's phone number to be "505-526-5466", or I may simply not know what you believe is his phone number.

If A knows what B believes is the value of P, then A may decide whether B is correct or incorrect in his belief, according to what A believes is the correct value, i.e., if I believe that John's phone number is 505-525-5466, and I believe that you believe it's 505-526-5466, then 1 believe that you are incorrect in your belief about John's phone number.

If A does not know what B believes is the value of P, it is possible that A can use other information to determine whether B is correct or incorrect in his belief about the value of P. For example, if you have not seen John in years, and he has moved house since you last saw him (but you do not know that he has moved house) although you believe that you know his phone number, I believe that you have his phone number wrong. I can believe this even if I do not know what his previous phone number was.

The ability to represent meta-beliefs, of the type described above, is crucial to. the process of ascribing belief. Lambda Formulas are an extension of lambda expressions that allow the meta-beliefs of a belief to be represented with their associated belief.

4.3. Lambda Formulas

Lambda expressions are finite partial functions. They are computable and map into a finite range of values. We extend the lambda expression representar tion to a Lambda Formula so we can represent opposing views on the value of such an expression[4]. For example,

(PHONE- NO- OF John) Be
{\(z){PHONE- NO- OF x) John) < John >

represents an expression that evaluates to John's phone number for John, but is unevaluable to any other agent. The imposition of John as a capable evaluator of the lambda expression (henceforth X-expression) is equivalent to adding an extra parameter to the expression; the extra parameter being the agent performing the evaluation.

The list of capable evaluators is not necessarily a constant because our beliefs may vary regarding who are the capable evaluators of an expression. Also different evaluators may evaluate the expression to different values, that is also subject to each individual's beliefs.

We want to express different agents as being capable of evaluating such an expression but to different values. We propose a more complex form of the capable evaluators list proposed in (Wilks, 1986). A typical expression in the augmented representation is shown in (figure 9).

$$\begin{array}{l}(\mathit{CURE\text{-}FOR\ Tuberculosis})\ \mathit{BE}\\ \quad (\lambda(x)\ (\mathit{CURE\text{-}FOR\text{-}TB}\ x))\\ \qquad \left[\begin{array}{c} <\{\mathit{MDs},\ \mathit{John}\},\ v_0>,\\ <\{\mathit{Dan},\ \mathit{Tibetan\text{-}Priest}\},\ v_1>,\\ \vdots\\ <\{\mathit{Sally},\ \mathit{Paul}\},\ v_n>,\end{array}\right]\end{array}$$

Figure 9. Augmented Lambda Expression for the Cure for Tuberculosis

Figure 9 contains n sets of capable evaluators and the values to which they evaluate the expression. The basic set of capable evaluators has been replaced by a set of pairs, consisting of a set of capable evaluators and the value that they return. A typical entry, for example, <{Sally, Paul}, v„>, is called a

Capable JEvaluators-Value pair (abbreviated to CEV pair). It is stipulated that each value «, must be of the same structure and that the X-expression returns this structure.

The set of CEV pairs is a function table representation of the X-expression. We refer to the X-expression and its function table as a X-formula.

4.4. Unknown Values

In the original representation, values that are unknown to an agent are represented by the agent not being a capable evaluator. In other words, awareness of the value of a X-expression is represented by being able to evaluate the X-expression. In figure 8 this would be equivalent to replacing the section below *"aware of the value of P"* by *"the correct value."* The original representation does not handle problems of awareness of P with sufficient power.

Being capable of evaluating the expression is still the criterion for knowing what the expression represents. Now, however, by the introduction of what we call *Uncertain Values* and *Uncertain_but_Jncorrect Values,* we are able to express more with the representation.

An uncertain value is a value (of a CEV pair) that is unknown to an agent, i.e., the agent is aware that the value exists but not what the value is. This uncertain value may or may not be equal to a known value or some other uncertain value. Thus, the agent can represent the value that another agent believes is the value of a X-expression without knowing what the value is.

Uncertain_biit_Jncorrect values are known by the agent to be different from the value that the agent believes to be correct. In other words the value, while unknown to the agent, is believed by the agent to be wrong (the meta-beliefs that A believes that what B believes is the value of P is wrong, even though A does not know what B believes is the value of P).

We have demonstrated representations of all cases where agent B (figure 8) is believed to be aware of P. Next, we consider the case where B is unaware of P.

4.5. Awareness, and Explicit & Implicit Mention of Agents

A question is *"how to represent an agent being aware of the proposition represented by a \-formula?"* Our solution is this: An agent is said to be aware of the proposition represented by a X-formula if and only if that agent is either explicitly or implicitly mentioned in the set of all agents that appear in the table for the formula; otherwise the agent is said to be unaware of the proposition.

An agent is explicitly mentioned if the agent is a member of one of the sets in the table and is implicitly mentioned if the agent is a member of a class of agents[5] that is explicitly mentioned in the table. This may be seen in example 5 where MDs are capable evaluators. Hence, any agent who is an MD is capable of evaluating example 5. MDs are explicitly mentioned in this example, someone who is an MD is implicitly mentioned.

An agent who is neither explicitly nor implicitly mentioned is deemed to be unaware of the proposition that the formula represents. In the case where an agent is implicitly mentioned in more than one class of agent, the most specific class is chosen to determine that agent's view.

4.6. Function Transformations

The use of X-formulas poses a question, what affect does ascribing a X-formula to an agent have upon the X-formula?

Ascribing a X-formula often requires altering the formula. For example, ascribing a formula to an agent, who does not know the value it returns for any agent, will require changing the formula to reflect this situation. Ascribing a X-formula to an agent thus involves changing the function table for the formula. This is equivalent to saying that ascribing a X-formula to an agent involves changing the function that the formula represents.

Function Transformations on X-formulas are discussed in more detail in (Ballim, 1987).

4.7. X-expressions and Meta-Beliefs

We can see how X-formulas solve the problems that were discussed in section 4. Lack of awareness is realised within X-formulas by an agent being neither an explicit nor implicit agent of the formula's capable evaluators set. Atypical beliefs are characterised by their associated meta-beliefs. X-formulas are very effective at representing meta-beliefs of atypical beliefs.

5. A Description of *ViewGen* and a Medical Example

VicwGen is a Prolog program that generates nested points of view. The program has a set of beliefs that are considered held by an agent known as the system. These beliefs are partitioned into topics (held in labelled PTs). An example call to the program is shown in (figure 10).

```
|?- viewgen( [system],[thalassemia] ,_J.

Viewpoint of [system]
 is
[
 thalassemia  isjtype_pf
    lambda(type_pf(thalassemia),genetic_disorder)
```

Figure 10. The System's View about Thalassemia[6]

The point of view represented in (figure 10) is the system's view about thalassemia. This view consists of one belief, namely that it is a genetic disorder.

ViewGcn operates by forming a contextual focus from the PTs of the terms in the list that forms the second argument to the program. The first argument is a list of agents. This list represents the nesting that is required. The third argument is a variable that is instantiated to the structure that the program returns.

In the following example assume that the system is a medical expert conducting a dialogue with a married couple who are seeking advice about thalassemia. The system has a lambda formula for the intensional description "*type_of(thalassemia)*" shown in (figure 11).

```
lambda(type_of(thalassemia), D,
            [
             [ [system, medicallyJnformed_person],
                                genetic_disorder ],
             [ [avg_man], disease ]
            ))
```

Figure 11. Lambda Formula for the type of thing that Thalassemia is

With the information that thalassemia is a type of genetic disorder the system is able to generate what it believes about them both. This point of view is shown in (figure 12).

```
|?-  viewgen([system],[thalasscmia,genetic_disorder],__).

Viewpoint of [system]  about
      [thalassemia,genetic_disorder]   is
(
  (  (D  is__type_of genetic_disorder and
      (Al suffers_from D  and
        (A2 suffersJYom D  and
          (Al  married_to A2 and A3 child_of [A1,A2]))))
   implies A3 suffersJYom D  )

  thalassemia is_type_pf
       lambda( type__of( thalassemia) ,genetic_disorder)
)

yes
|?-
```

Figure 12. The System's view of Thalassemia and Genetic__disorder(s)

The point of view in (figure 12) contains the belief that thalassemia is a genetic disorder. In addition it contains a complex belief the child of two people, who both have a genetic disorder, will suffer from the same genetic disorder from them[7].

Suppose that the couple in question are called Paul and Sally, and that the system has the beliefs about them shown in (figure 13).

```
|? - viewgen( [system],[sally,paul], J.

Viewpoint of [system]  about [sally,paul]
 is
|
  sally is_type__of medically_jnformed_person
  sally married_to paul
  paul suffersJYom thalassemia
  sally suffersJYom thalassemia


yes
|*.
```

Figure IS. The System's view of Sally and Paul

The important beliefs here are that the system believes that Sally is a medically informed person, and that they both suffer from thalassemia. We presume that we have reasoning mechanisms that allow us to reason with the viewpoints in figures 13 and 14 to hypothesise that any offspring of the pair will also suffer from thalassemia (something that they need to be warned about). If we presume that the system follows Gricean Maxims then we need to generate what the system believes each of them believes about thalassemia. These points of view are shown in figures 15 and 16.

```
|?- viewgen( [system,sally], [thalassemia],_).

Viewpoint of [system,sally]  about [thalassemia]
 is
[
  thalassemia is_type_pf
             lambda( type_of( thalassemia), genetic_disorder)
i

yes
|?-
```

Figure 14. The System's View of Sally's View of Thalassemia

```
|?-  viewgen([system,paul],[thalassemia],^).

Viewpoint of [system,paul]  about [thalassemia]
 is
[
  thalassemia is_type_of
             lambda( typejof (thalassemia), disease)


yes
|?-
```

Figure 15. The System's View of Paul's View of Thalassemia

In (figure 14) the belief that thalassemia is a type of genetic disorder has been ascribed to Sally. This is because the system believes that Sally is a medically informed person, and the system believes that medically informed people know that thalassemia is a genetic disorder (see figure 11). Paul, however, is not believed to be a medically informed person. He is simply an average man[8]. The average man view of thalassemia is that it is a disease. This is just the view that the system believes Paul has about thalassemia (figure 15).

In a dialogue between three agents it can be necessary for one agent to model what the second believes the third believes, or what the third believes the second believes about the topic. So we generate what the system believes Sally believes Paul believes about thalassemia, and what the system believes Paul believes Sally believes about thalassemia.

Using the beliefs about Sally and Paul the system generates that Sally believes that Paul believes that thalassemia is a disease. However, the system has generated that Paul believes that Sally believes thalassemia is some kind of complex thing that he doesn't know about, represented by ^{t4}[u,l]" which is an uncertain value. In other words, according to the default rule Paul believes that Sally is a medically informed person (since there is no evidence to contradict this) and hence Paul believes that Sally has a more complex belief about thalassemia than he does.

6. Conclusions

This work claims only to be a first implementation of a "belief engine" (Maida, 1986) that contains plausible heuristics or th« default ascription of concrete, contentful beliefs. Such a process will be needed by any project that proposes (e.g. Pollack, 1986) to model the interaction of agents planning on the basis of differing beliefs and plans. We believe no other system has yet tackled this problem in a practical way. Many extensions will be required to the current work, particularly in the treatment of the identication of intensionally distinct but extensionally identical individuals and classes.

Reference*

Attardi, G. and Simi, M. (1984) Metalanguage and Reasoning Across Viewpointe. Proc. of ECAI-84, 315-324.

Ballim, A. (1986) Generating Points of View. Memoranda in Computer and. Cognitive Science, MCCS-86-68, Computing Research Laboratory, New Mexico State University, Las Cruces, NM 88003, USA.

Ballim, A. (1987) The Subjective Ascription of Belief to Agente. In AISB-87 Conference Proceedings, J. Hallam & C. Mellish (eds), John Wiley & Sons: Chichester, England.

Barn den, J. (1983) Intensions as such: an Outline. In Proceedings of the 8th International Joint Conference on Artificial Intelligence, 280-286.

Creary, L. (1979) Propositional Attitudes: Fregean Representation and Simulative Reasoning. In Proceedings of IJCAI-7Q, Tokyo.

Konolige, K. (1985) A Computational Theory of Belief Introspection. In Proceedings of IJCAI-85, Los Angeles, 502-508.

Maida, A. S. (1986) Introspection and reasoning about the Beliefs of other Agents. In Proceedings of Cognitive Science Society, 187-195.

Minsky, M. (1975) A Framework for Representing Knowledge. In The Psychology of Computer Vision, Winston, P.H. (ed.). McGraw-Hill, New York.

Moore, R. and Hendrix, G. (1979) Computational Models of Belief and the Semantics of Belief Sentences. SRI Technical Note No. 187.

Perrault, R. k Allen, J. (1980) A plan-based analysis of indirect speech acts. In Amer. Jnl. of Computational Linguistics, 6, 167-182.

Pollack, M.E. (1986) A Model of Plan Inference that Distinguishes Between the beliefs of Actors and Observers. In Proceedings of the Association for Computational Linguistics, 207-214.

Shadbolt, N. (1983) Processing Reference. In Jnl. of Semantics, Vol. 2, no. 1, 63-98.

Shapiro, S.C. and Rapaport, W.J. (1986) SNePS Considered as a Fully Intensional Propositional Semantic Network. In Proceedings of AAAI86, 278-283 .

Wilks, Y.A. (1977) Making Preferences more Active. Artificial Intelligence, Vol. 8, 75-97.

Wilks, Y.A. (1986) Relevance and Beliefs, In Reasoning and Discourse Processes, Brown, K., Myers, T., and McGonigle, B., (eds.). Academic Press, New York.

Wilks, YA. and Bien, J. (1979) Speech Acts and Multiple Environments. In Proceedings of IJCAI-79, Tokyo, 451-455.

Wilks, Y.A. and Bien, J. (1983) Beliefs, Points of View and Multiple Environments. In Cognitive Science 8, 120-116.

Notes

1) In Figure 3 the environment in question is that portion that is inside the section marked "Frank".

2) Two distinct forms of pushing down environments are identified in (Ballim 1986) only one of which is discussed here.

3) The other agent may be A himself. In that case the meta-belief is ctae of introspection, however it would be absurd for agent A to have a belief that he is unaware of belief P so this situation is not allowed. Introspection is discussed in detail in Konolige(l985) and in Maida(1986).

4) It may appear that talking about the value of a proposition is only valid for knowledge of values, such as phone number, etc., but not for propositions of the form "John is here", however, these propositions may be treated in one of two ways: (a) they may often be rewritten to reflect a value (e.g., location of John is x), or (b) the truth or falsity of the proposition may be treated as its value. Coupling the latter case with unknown values (see section 4.4) provides us with the power of a Kleene three-valued logic, or a modal logic with a " not-know-whether" operator.

5) Meaning Postulates are used to determine if an agent is a member of a particular class of agents.

6) In the examples that we use beliefs are simply propositions that reside within a point of view. This propositions are skolemised so that existentially quantified variables are replaced by skolem constants and all other variables are universally quantified within a viewpoint. Variables are terms that begin with a capital letter.

7) This is an extremely simplistic belief, but suffices for the purpose of this example. A real medical expert using VicwGcn would contain far more complex beliefs than this.

8) VicwGcn has a lattice of classes and a set of meaning postulates that allow Paul to be seen as an average man.