Consiglio Nazionale delle Ricerche



Istituto di Elaborazione dell'Informazione

# De Dicto Reading of Predicate Symbols in Artificial Intelligence

Alessandro Saffiotti & Fabrizio Sebastiani

Nota Interna IEI-B4-39-1990

Via S. Maria, 46 - 56126 Pisa (Italy)
Phone : +39.(0)50.553159
Fax : +39.(0)50.554342
Telex : 590305 IEICNR I

# *De Dicto* Reading of Predicate Symbols in Artificial Intelligence

Alessandro Saffiotti
Université Libre de Bruxelles
IRIDIA
50 Avenue F.D. Roosevelt, CP 194/6 - 1050 Bruxelles (Belgium)

Fabrizio Sebastiani
Istituto di Elaborazione dell'Informazione
Consiglio Nazionale delle Ricerche
Via S. Maria, 46 - 56126 Pisa (Italy)

Technical Report IEI-B4-39-1990

September 1990

## Abstract

After an introductory discussion on the role of logics for mental attitudes in computational models of dialogue, we present the M-KRYPTON representation language. M-KRYPTON has been built in order to provide a tool for reasoning about attitude reports in the presence of multiple cognitive agents, and where an attitude report is interpreted *de dicto*, i.e. relative to (the reporter's view of) the cognitive space of the agent the attitude is attributed to. M-KRYPTON may be seen as an attempt to combine insights from AI knowledge representation and doxastic modal-like logics. On one side, hybrid knowledge representation systems embody a powerful representational paradigm, accounting for both belief (or, popularly, "knowledge") about the "necessary" nature of concepts relevant to the domain of discourse and belief in "contingent" facts about them. On the other side, the possible worlds semantics typical of modal logic is easily tailored to model various notions of a doxastic nature that are relevant to the modelling of cognitive agents. By recasting the KRYPTON hybrid KR system in terms of possible worlds semantics we have obtained a semantic account that, besides being "functionally" equivalent to the original one, is easily extensible to deal with operators for mental attitudes. In particular, we have concentrated on adding to KRYPTON the possibility of representing beliefs about the beliefs of multiple agents, where such agents may believe in propositions either of a "necessary" or of a "contingent" nature.

## 1 Introduction

The work that will be described in this paper is being carried out in the frame of a project that attempts to create an adequate computational model of "extended" man-machine dialogue in natural language that also incorporates techniques for the detection and repair of miscommunications. By specifying that the kind of dialogue that is to be dealt with is "extended", we mean that the model must account for features of dialogue (such as mixed initiative, indirect speech acts, ellipses, anaphora, dependence from the pragmatic context, etc.) that, although are seldom dealt with in existing natural language man-machine interfaces, are so central to dialogic communication to be considered essential to any model of dialogue that claims to have this name.

Borrowing from the terminology of computational complexity theory, natural language processing has sometimes been informally referred to as an "AI-complete problem", meaning that it includes *in nuce* the most challenging and conceptually difficult problems today facing Artificial Intelligence. Dialogue processing is perhaps the most difficult task among the ones related to natural language, and one for which no strong guidelines to the development of computational models do exist.

The reason of this inherent difficulty is perhaps to be found in the heavy dependence of dialogic communication on (both linguistic and pragmatic) context[1]. Following the terminology introduced by Barwise and Perry (1983), we might say that dialogic communication is *efficient* in nature; that is, context dependence is so far-reaching a phenomenon in dialogue to relegate context independent utterances (utterances realized by what Quine (1960) calls *eternal sentences*) to the status of mere exceptions.

## 2 Dialogue processing and mental attitudes

Dialogue diverges significantly from other uses of language (e.g. written text) in its heavy use of the pragmatic context in which the dialogue takes place. For instance, in the following dialogue (exemplifying use of an "indirect speech act" (Searle 1975)):

A : Do you know the time?
B : Yes. It's four o'clock.

B's response is influenced by B's belief that A does not really want to know whether B knows what the time is, but wants actually to know the time. In general, speakers may be said to "plan" their contributions to the dialogue they are participating in, and to take into account the partner in the conversation and her cognitive reality as one of the fundamental parameters influencing their

1 See (Hirst 1981) for a stimulating account of the wide variety of influences of context on language.

planning activity.

In computational models of dialogue, dependence from the pragmatic context is enforced by means of an explicit representation of all the features of context that are relevant for the task of dialogue participation, and on the subsequent exploitation of these representation structures during the production/assimilation of utterances. For this purpose, AI knowledge representation languages are employed[2]. Unfortunately, knowledge representation lends too weak a handle on the problem, in that typical KR languages are geared towards the representation of an "inanimate" reality, of a world where no entities endowed with a cognitive activity are present. More precisely: although by making use of such languages we may represent sentences that do involve an "animate entity", as

John is a friend of Bill

we are not able to represent sentences involving his/her cognitive activity, as

John wants Bill to be his friend

This is due to the fact that typical KR languages (see e.g. (Vilain 1985) or (Brachman *et al.* 1985)) are *extensional* in nature, while notions having to do with a cognitive activity (e.g. believing, knowing, seeing, wanting, intending, etc. - the so-called *mental attitudes*) have an *intensional* nature[3]. Perhaps the most significant way in which mental attitudes show their intensional nature is their unusual behaviour with respect to Leibnitz's principle of substitutivity, according to which the terms of a true statement of identity are everywhere intersubstitutive, *salva veritate*. That is, while from

(1) The Earth moves
(2) The Earth = the third planet from the Sun

we may infer

(3) The third planet from the Sun moves

<u>it is not obvious that</u> from

(4) Galileo believes that the Earth moves
(2) The Earth = the third planet from the Sun

we may infer

---

2 See (Brachman & Levesque 1985) for a representative collection of papers on knowledge representation.

3 See (Carnap 1947) for the original formulation of the extension/intension distinction.

(5) Galileo believes that the third planet from the Sun moves

Whether this inference is reasonable or not depends on the interpretation we give to the belief operator. Under one interpretation (the so-called *de dicto* reading) the "believes" operator applies to the whole proposition (e.g. the one denoted by "the third planet from the Sun moves"); according to this reading, the inference in our example is not legitimate, since Galileo might be unaware that the Earth is actually the third planet from the Sun. Under another interpretation (the so-called *de re* reading) the "believes" operator is part of a complex predicate that applies to a certain entity; according to this reading, the inference in our example is legitimate, since (5) is interpreted as saying something like "Galileo believes of the third planet from the Sun that it moves". In other words, in the *de dicto* reading the noun phrase "the third planet from the Sun" is interpreted idiosyncratically to the agent to whom the attitude is attributed to, while in the *de re* reading it is interpreted idiosyncratically to the reporter of the attitude[4].

As the *de dicto* reading is somehow the standard reading of attitude reports (or, at least, is the more evident to most people), dealing with them requires logics and languages that do not license unrestricted substitutivity of equals. One of the most popular approaches to this problem, both in the philosophy of language and, more recently, in AI, has been the application of the formal apparatus developed for modal logic to the modelling of mental attitude notions. This approach, pioneered by Hintikka (1962, 1969), is a natural approach to the problem, in that 1) the modal notions of necessity and possibility also give rise to "opaque" contexts (i.e. contexts where substitutivity of equals is not allowed); 2) there is a substantial and well-established body of work on modal logic[5]; 3) the standard semantics for modal logic is easily tailored to deal with many interesting notions relevant to the modelling of cognitive agents.

## 3 Inadequacies of modal-like logics for mental attitudes

Logics for mental attitudes that are built by relying on the formal apparatus of modal logic, although perhaps the most viable means for modelling attitudes developed to date, suffer from a series of problems. We might divide these problems in roughly two groups:

1) problems related to expressive adequacy

4 Although philosophers have concentrated on spanning the *de dicto/de re* distinction as applied to singular terms, it is patently obvious that a similar distinction also applies to other non-logical symbols (e.g. predicate symbols). As remarked by Bonomi (1983:185), these problems are no different, in that

"the distinction has to do with reports of mental events in general, and is grounded on the possiblity to situate in alternative conceptual spaces the interpretation of the expressions involved"

(our translation). This issue is taken at heart in M-KRYPTON.

5 See (Hughes & Cresswell 1968) for an extensive introduction to modal logic, mainly focusing on syntactic aspects; see (Chellas 1980) for a more semantically-oriented introduction.

2) problems related to computational complexity

The work we describe in this paper may be seen as an attempt to give an answer, albeit limited, to both these kinds of problems.

## 3.1 Expressive adequacy and awareness

Perhaps the most compelling problem having to do with expressive adequacy is *consequential closure* (better known, in the case of knowledge and belief, as *logical omniscience* (Hintikka 1975))[6]: an agent is modelled as believing in all logical consequences of what she already believes. Needless to say, the fact that real agents are not logically omniscient renders this approach highly unsuitable for applications where such agents are to be modelled.

In a recent paper, Fagin and Halpern (1988) have argued that this substantial inadequacy arises from a number of different features that characterize the belief of real agents: lack of awareness, resource-boundedness and incomplete focussing. In particular, according to the authors, lack of awareness has to do with the fact that we cannot even consider saying that an agent believes that p or that ¬p when p is a proposition she is not even aware of. To quote an example from their paper, lack of awareness explains the fact that a Bantu tribesman possibly neither can believe that

(6) The price of personal computers is going down

nor can believe in its negation.

Fagin and Halpern indicate a possible solution ("the logic of general awareness") to this problem; roughly, it is based on the idea that there is a set of propositions that an agent is aware of, and that ("explicit") belief may be seen as classical ("implicit") belief modulo the awareness set. In other words, an agent can explicitly believe that p only if she is aware of p; hence, unlike in classical doxastic logics, the belief set of an agent is not necessarily closed under logical consequence. However, Fagin and Halpern admittedly dodge the problem of investigating the nature of awareness, and propose an essentially syntactic operator that allows to tailor the notion of awareness to be captured to the specific application in mind. In practice, they are interested in providing <u>useful</u> (rather than cognitively plausible) "filters" for explicit belief, leaving untouched the question whether these filters correspond to our intuitive notion of "awareness" (and, as a consequence, what is explicit belief).

Investigating this notion more deeply is one of our primary interests, as we are convinced that

6 From this point on we will no more speak of logics for mental attitudes in general, concentrating instead on logics for belief ("doxastic logics"). This is not a limitation, as most other attitudes share the same characteristics of belief, and can therefore be modelled in similar ways.

a cognitively plausible notion of awareness is potentially relevant to formal accounts not only of belief, but also of all other attitudes. To this respect, we think that in the same situation our Bantu tribesman also neither can want that (6) hold, nor that its negation does.

The notion of awareness we propose (Saffiotti & Sebastiani 1989) relies on the intuitive idea that awareness of a proposition p be somehow identifiable with the "capability of attributing a meaning to p". In this light, a closer inspection of the Bantu tribesman example suggests that unawareness might be traced back to deeper roots than those captured by the analysis of Fagin and Halpern. The Bantu tribesman might be unaware of the cited proposition because of unacquaintance with the concept "personal computer", or with the concept "price"; alternatively, he might be acquainted with all of these concepts and still be unaware of the proposition because he might be at odds with the connection between the concepts that the proposition seems to hint at (e.g. he might believe that personal computers are priceless objects). A common trait shows up throughout the situations we have depicted: the source of unawareness can be traced back to the Bantu tribesman's "conceptual vocabulary", to his "terminology".

The strongly terminologic nature of awareness calls for a doxastic logic where the impact of the agent's terminology on the proposition ("the terminologic content of the proposition") can be singled out, hence allowing to characterize the fact that an agent is aware of a proposition in terms of this component alone.

## 3.2 Computational complexity and awareness

The most important drawback hindering the utilization of modal-like doxastic logics in practical AI systems is that they are computationally intractable. For instance, it is well known (Ladner 1977; Halpern & Moses 1985) that the problem of deciding satisfiability in the propositional modal logic K is PSPACE-complete. As one might expect, if one adds other features to the logic (such as multiple agents, common belief and/or group belief, quantifiers), the situation grows even worse; nor reverting to different modal logics (such as *T*, *S4*, *S5* or *KD45*) is of any help. As the importance of applying such logics seems to be out of question, it is of primary importance to look for ways of getting around the problem and limiting the computational load of such reasoning tasks.

In this respect, the above-mentioned notion of "awareness of a proposition" should prove of keystone importance: a formula involving belief in a proposition can be shown false also by showing that the agent is unaware of the proposition, hence "bypassing" the normal proof procedure.

Of course, this does not have any impact on the computational complexity of reasoning about belief, as it is not possible to quantify the percentage of false formulae that are such on grounds of awareness: this obviously depends on the composition of the set of formulae that are fed to the theorem prover. The incidence of awareness might thus be characterized at best statistically, and the notion of statistics is outside the realm of computational complexity theory. What is actually argued here is not that the complexity of reasoning gets diminished, but that some cases are singled out in which the invocation of the reasoning process may be avoided.

Of course, if awareness has to ease computation of attitude formulae, it has to be easily computable itself. To this respect, enforcing a terminology-based notion of awareness, such as the one we have developed, is promising: in fact, Brachman and Levesque (1984) have shown that reasoning about terminology can be done in polynomial time if the representation language is carefully designed.

## 3.3 A language for terminologic and assertoric beliefs

A basic requirement for modelling a terminology-based notion of awareness is the availability of a language that allows both "assertoric" formulae expressing factual truths (as, for instance, statement (6)) and "terminologic" formulae expressing intended relationships among the concepts referred to in assertoric formulae (as, for instance, "A thatched house is a house whose roof is made of straw"; see (Woods 1975) and (Brachman & Levesque 1982) for more on the distinction between assertoricity and terminology). The M-KRYPYON representation language, that we will discuss from section 5 onwards, was built with these goals in mind.

Remarkably, the framework we have developed accounts both for the substantial difference and for the intended relationship between terminology and assertoricity, with the former constraining the latter but not vice-versa. The marriage between notions typical of knowledge representation and ones developed in the field of doxastic logic turns out to be a distinctive characteristic of our logic.

This paper describes our attempt to meet the representational needs of our model of awareness by extending the KRYPTON hybrid knowledge representation system. This attempt has resulted in a knowledge representation language (we have called it M-KRYPTON) that, though designed with our application in mind, has a character of generality, and is particularly geared to modelling situations where multiple reasoning agents are involved (such as user modelling (see e.g. Wahlster & Kobsa 1986) or communication protocols analysis (see e.g. Halpern & Moses 1984)). After a brief overview of KRYPTON (section 4) we give an alternative semantics for the KRYPTON assertoric and terminologic languages which seems to provide a more natural environment for developing the desired extensions; we also state its "functional" equivalence to the semantics originally given by Brachman and colleagues (section 5). We then go on to describe our own extensions (section 6); these include developing a more unified approach towards assertoricity and terminology, plus dealing with the representation of beliefs about the beliefs of multiple agents.

# 4 An overview of KRYPTON

KRYPTON is a hybrid knowledge representation system that comprises two separate representation languages:

1) a terminologic language, by means of which concepts and functional roles of these concepts can be described; for instance, the concept of a Logician can be expressed in KRYPTON by means of the term (<u>VRGeneric</u> Scientist Subject Logic), where Scientist and Logic are themselves names of concepts and Subject is a name of a functional role;
2) an assertoric language, by means of which facts about the domain of discourse may be described that possibly make reference to such concepts and roles; for instance, the fact that some logicians are also poets can be expressed in KRYPTON by the formula (EXISTS x (AND (Logician x) (Poet x))), where Logician and Poet are names of concepts[7].

Both the differences and the relationships between the two languages are accounted for by means of what the authors call a "hybrid semantics".

The other key feature of the KRYPTON system is that it embodies what Levesque (1984b) calls a "functional" approach to knowledge representation. This means that a user interacts with a KB (seen as an abstract data type) only via an interface language consisting of a small set of precisely defined operations; thus, she may not be (and should not be) concerned with the actual data structures and mechanisms involved.

In this section we first give a complete description of the two representation languages together with an informal sketch of the underlying semantics; subsequently, we go on to briefly describe the interface-level operations that embody the "functional" specification of KRYPTON KBs. Space precludes us from giving a detailed formal description of the semantic account originally given to KRYPTON by its authors; for this we refer the reader to (Brachman *et al.* 1985).

The assertoric language of KRYPTON is a pure (i.e. with no individual constants and function symbols) first-order language. Hence, an assertoric well-formed formula is defined according to the following syntax:

<wff> ::= (<n-ary predicate symbol> <var>$_1$ ... <var>$_n$) | $n \geq 0$
(NOT <wff>) | (OR <wff> <wff>) | (EXISTS <var> <wff>)

In addition, the metalinguistic abbreviations AND, IMPLIES and FORALL are defined in the obvious way.

The terminologic language permits the formation of complex terms ("gterms") according to the following syntax:

7 The distinction between terminology and assertoricity recalls some dualisms that are well known in the philosophy of language: analytic/synthetic, necessary/contingent, *a priori/a posteriori*. It is not clear to which notions assertoricity and terminology might be analogous, the reasons being that 1) these notions themselves do not have universally accepted definitions; 2) in different papers the proponents of the assertoricy/terminology distinction seem to suggest different notions (compare (Brachman & Levesque 1982; Levesque 1984b; Israel and Brachman 1984; Brachman et al. 1985)); 3) in only one paper (Brachman et al. 1985) they provide a formal semantic account; according to this work, the similarity of assertoricity/terminology to necessity/contingency is apparent.

| | | | |
|---|---|---|---|
| <gterm> | ::= | <n-ary predicate symbol> \| <concept> \| <role> | $n \geq 0$ |
| <concept> | ::= | <unary predicate symbol> \| | |
| | | (ConGeneric <concept>$_1$ ... <concept>$_n$) \| | $n \geq 0$ |
| | | (VRGeneric <concept>$_1$ <role> <concept>$_2$) | |
| <role> | ::= | <binary predicate symbol> \| | |
| | | (RoleChain <role>$_1$ ... <role>$_n$) | $n \geq 0$ |

We can thus express concepts obtained by role value restriction, such as the above mentioned Logician concept, by conceptual conjunction, such as (ConGeneric Logician Woman) ("a woman logician"), as well as roles obtained by role composition, such as (RoleChain Parent Sister) ("a sister of a parent", i.e. an aunt).

Predicate symbols constitute the link between the two languages: each of them may be associated to a gterm which becomes its definition, as well as used in an assertoric wff.

As a starting point towards an integrated semantics for KRYPTON, the two languages are first given a semantics independently from each other: the *truth* of wffs is defined by mapping them into truth values, while the *extension* of gterms is defined by mapping them into sets of individuals.

Bridging the gap between the two languages starts from the notion of *symbol table*, i.e. a function S that associates to each predicate symbol of the assertoric language a gterm, which becomes its definition. Consideration is then restricted to those notions of truth and extension that are "consistent" with the contents of a given symbol table S: these are called *truth valuations* wrt S and *extension functions* wrt S, respectively.

All the extension functions consistent with a symbol table S are then considered together in the definition of a *subsumption* relationship between gterms: a gterm $g_1$ is said to subsume gterm $g_2$ (wrt S) iff its extension is a superset of $g_2$ according to every extension function wrt S. Abstracting away from particular extension functions results in the required relationship between terminology and assertoricity, with the former constraining the latter but not vice-versa. An ordered pair formed by a subsumption relationship $\Rightarrow$ and a truth valuation $\omega$ is called an *outcome*, and may be seen as a complete specification of a state of affairs and of a "conceptual dictionary".

Next, we turn to briefly describe the "functional" level of KRYPTON. The fact that a KB is usually an incomplete specification of a state of affairs and of a conceptual dictionary is captured by defining a KRYPTON KB as the set of all outcomes that are consistent with this specification. The following clauses define the main functional-level operations of KRYPTON:

| | | |
|---|---|---|
| NEWKB [] | = | $\{\langle\Rightarrow, \omega\rangle \mid \langle\Rightarrow, \omega\rangle$ is an outcome$\}$ |
| ASK [$\alpha$, KB] | = | yes if for each $\langle\Rightarrow, \omega\rangle$ in KB, $\omega(\alpha)$ = true; no otherwise |
| TELL [$\alpha$, KB] | = | $\{\langle\Rightarrow, \omega\rangle$ in KB $\mid \omega(\alpha)$ = true$\}$ |

DEFINE [t, g, KB] = $\{\langle\Rightarrow, \omega\rangle$ in KB $\mid t \Rightarrow g$ and $g \Rightarrow t\}$
SUBSUMES [$g_1$, $g_2$, KB] = yes if for each $\langle\Rightarrow, \omega\rangle$ in KB, $g_2 \Rightarrow g_1$;
no otherwise

where $\alpha$ ranges on assertoric wffs, g, $g_1$, and $g_2$ on gterms, t on unary and binary predicate symbols and KB on KRYPTON knowledge bases. NEWKB creates an empty KB, TELL and DEFINE update a KB and ASK and SUBSUMES query it; roughly speaking, the semantic counterpart of an update operation is the restriction of the number of outcomes of the KB, while the one of a query operation is their inspection.

We close this section with some informal comments on the properties of these operations as deriving from the semantics given in (Brachman *et al.* 1985) and sketched above.

First of all, it might be the case that ASK [$\alpha$, KB] = no and, at the same time, ASK [(NOT $\alpha$), KB] = no.

Next, each valid formula of the first order predicate calculus holds in all outcomes (i.e. in NEWKB []). Analogously, in all outcomes the relation of subsumption comprises all the expected pairs of terms (e.g. $g_1$ subsumes (<u>ConGeneric</u> $g_1$ $g_2$)), as resulting from the semantics of the above described term calculus.

Third, properties that seem reasonable to expect of the interaction between ASK and TELL and between SUBSUMES and DEFINE do in fact hold for any KB, such as

ASK [$\alpha$, TELL [$\alpha$, KB]] = yes
SUBSUMES [t, g, DEFINE [t, g, KB]] = yes
SUBSUMES [g, t, DEFINE [t, g, KB]] = yes

Finally, also the interaction between terminology- and assertoricity-related operations works well: for instance, for every KB

ASK [(FORALL x (IMPLIES (P x) (Q x))), DEFINE [P, (<u>ConGeneric</u> Q R), KB]] = yes

while it is not true in general that

SUBSUMES [P, Q, TELL [(FORALL x (IMPLIES (P x) (Q x))), KB]] = yes

## 5 An alternative semantics for KRYPTON

In this paragraph we propose an alternative semantic account of the KRYPTON system in the style of the "possible worlds semantics" (PWS) originally developed for modal logic. The reasons for undertaking such a task are manifold.

First of all, the way KRYPTON distinguishes between terminology and assertoricity can be recast in terms of the distinction between necessary and contingent truths respectively, as enforced in classical modal systems.

Second, a PWS for KRYPTON appears more simple and intuitive than the "hybrid semantics" given in (Brachman *et al.* 1985), allowing, among other things, to highlight the difference between terminology and assertoricity by relating in a clearer way terminologic facts to truth in a multiplicity of states of affairs, and assertoric facts to truth in a particular state of affairs only.

Third, the possible worlds approach our semantics relies on is somehow more orthodox; in particular, a significant body of work on logics for mental attitudes has been carried on in this framework. In the next chapter we will see how this fact allows us to extend KRYPTON in a natural way to deal with modal operators for representing beliefs about the beliefs of multiple agents. As required, our framework will make it possible to distinguish between beliefs in propositions of an assertoric nature and in propositions of a terminologic nature. Remarkably, the possible worlds framework would allow us, if desired, to add to the language other modal operators (e.g. in order to model other mental attitude notions) with the same naturalness.

Fourth, our semantics suggests a natural extension of the KRYPTON language to express terminologic definitions as wffs, so yielding, in a certain sense, a truly unified hybrid language. This modification brings with it a uniform treatment of belief, because the object of belief will then be a proposition, regardless of its assertoric or terminologic nature.

In this section we will concentrate on our alternative semantics for the KRYPTON language, while in the next one we will discuss our extensions.

Before specifying our semantics in detail it is perhaps worthwhile to briefly recall the basic notions of possible worlds semantics. PWS was formalized by Kripke (1963) in terms of what are now called *Kripke structures* to provide a semantic account of the modal notions of necessity and possibility. Roughly speaking, while standard Tarskian semantics works by mapping linguistic objects on a particular state of affairs ("interpretation"), PWS interprets these objects on a Kripke structure, which incorporates a number of alternative states of affairs ("possible worlds"), together with a binary relation (the "accessibility" relation) on these worlds. Intuitively, the accessibility relation tells us which worlds are possible alternatives given a particular one; accordingly, proposition p is necessarily true in world w if it is true in all worlds that stand in the accessibility relation with w.

Our PWS for KRYPTON is then given in terms of a Kripke structure $M=<S, D, V, \mathcal{T}>$, where S is a non-empty set of states, D is a domain of individuals, V is a function from individual variables to elements of D and from pairs $<p, s>$ (with p an n-ary predicate symbol and s an element of S) to n-ary relations over D, and $\mathcal{T}$ is a binary reflexive relation over S. A *possible world* is then defined as an ordered pair $w=<M, s>$; intuitively, a possible world represents a complete specification of a state of affairs.

The $\mathcal{T}$ relation allows us to capture the difference between assertoricity and terminology: intuitively, assertoric sentences have a "contingent" import, and their truth is thus evaluated with respect to a specific world, whereas gterms have a "necessary" import, and subsumption between them is thus evaluated with respect to a set of worlds (the ones accessible through $\mathcal{T}$).

Formally, we define the truth with respect to a world $<M, s>$ of an assertoric wff $\alpha$ of KRYPTON by means of a truth relation $\mid=$, where $<M, s> \mid= \alpha$ is to be read as "$\alpha$ is true in world $<M, s>$":

**Definition 1**: Let $<M, s>$ be a possible world, with $M=<S, D, V, \mathcal{T}>$ and s an element of S. A *truth relation* $\mid=$ is such that, for all predicate symbols p and for all wffs $\alpha$ and $\beta$:

| | | |
|---|---|---|
| $<M, s> \mid= (p\ x_1 \ldots x_n)$ | iff | $<V(x_1), \ldots, V(x_n)> \in V(p, s)$ |
| $<M, s> \mid= (NOT\ \alpha)$ | iff | it is not the case that $<M, s> \mid= \alpha$ |
| $<M, s> \mid= (OR\ \alpha\ \beta)$ | iff | $<M, s> \mid= \alpha$ or $<M, s> \mid= \beta$ |
| $<M, s> \mid= (EXISTS\ x\ \alpha)$ | iff | $<M_{[x/d]}, s> \mid= \alpha$ for some $d \in D$ |

where $M_{[x/d]}$ is defined as the Kripke structure that is exactly like M except that V(x)=d. A formula $\alpha$ is *valid* iff $<M, s> \mid= \alpha$ for all possible worlds $<M, s>$. Next, similarly to (Brachman *et al.* 1985), we have:

**Definition 2**: Let $<M, s>$ be a possible world, with $M = <S, D, V, \mathcal{T}>$ and s an element of S. Then, for any gterm g, the *extension* of g in world $<M, s>$ (written $E_{<M, s>}(g)$) is given by:

1. $E_{<M, s>}(p)$ is V(p, s) for any predicate symbol p.
2. $E_{<M, s>}$(<u>ConGeneric</u> $c_1 \ldots c_k$) is the intersection of the $E_{<M, s>}(c_i)$ for i = 1, ..., k, and D if k=0.
3. $E_{<M, s>}$(<u>VRGeneric</u> $c_1$ r $c_2$) is those elements x belonging to $E_{<M, s>}(c_1)$ such that $<x, y>$ is in $E_{<M, s>}(r)$ only when y is in $E_{<M, s>}(c_2)$.
4. $E_{<M, s>}$(<u>RoleChain</u> $r_1 \ldots r_k$) is the relational composition of the $E_{<M, s>}(r_i)$ for i = 1, ..., k.

Next, we define what it means, in terms of possible worlds, that gterms stand in a relation of subsumption:

**Definition 3**: Let $<M, s>$ be a possible world, with $M = <S, D, V, \mathcal{T}>$ and s an element of S. Let $g_1$ and $g_2$ be gterms. Then $g_1$ *subsumes* $g_2$ in $<M, s>$ iff for every t such that $s\mathcal{T}t$, $E_{<M, t>}(g_2) \subseteq E_{<M, t>}(g_1)$.

Hence the truth relation tells us which (assertoric) formulae are true in a given possible world $<M, s>$, while the subsumption relation tells us which (terminologic) relationships between gterms hold in a given possible world $<M, s>$. Such possible worlds play thus the same role outcomes play in the semantics given in (Brachman *et al.* 1985). As a matter of fact, each possible world supports the truth of all tautologies of the first order predicate calculus together with the logical consequences (in the sense of classical logic) of any other set of formulae whose truth it already supports. Moreover, each possible world "behaves well" with respect to definition 2: for instance, for any possible world and for any gterms $c_1$, r and $c_2$, (VRGeneric $c_1$ r $c_2$) is subsumed by $c_1$.

More interestingly, the reflexivity of the accessibility relation $\mathcal{T}$ guarantees the correct hierarchy between terminology and assertoricity by imposing that, for any pair of predicate symbols p and q, if p subsumes q in world <M, s>, then the formula (FORALL x (IMPLIES (P x) (Q x))) be true in <M, s>. In fact, if P subsumes Q in <M, s>, then the extension of P will be a superset of the extension of Q in any <M, t> such that t is accessible from s through $\mathcal{T}$ and in particular, by reflexivity, in <M, s> itself, so forcing the truth in <M, s> of the above formula. That the converse is not true in general is guaranteed by the fact that subsumption is evaluated with reference not only to world <M, s> but to all worlds <M, t> such that t is accessible from s through $\mathcal{T}$.

Similarly to the treatment given in (Brachman *et al.* 1985), we see a *knowledge base* as an incomplete specification of what the world and the conceptual dictionary are, and thus identify it with the set of all (completely specified) worlds which are consistent with it.

As for the "functional" level, the operations for the interaction with knowledge bases that we saw at the end of Section 2 can be redefined according to the new semantics, with possible worlds replacing outcomes. More precisely, it can be shown that our alternative semantics for KRYPTON has no impact on the functional level. To do this, we first need a definition (the * superscript identifies operations defined according to our semantics and KBs obtained through the application of these operations):

**Definition 4.** Let KB and KB* be two knowledge bases. We say that KB is *functionally equivalent* to KB* iff, for all wffs $\alpha$ and for all gterms $g_1$, $g_2$,

1) ASK [$\alpha$, KB] = ASK* [$\alpha$, KB*]
2) SUBSUMES [$g_1$, $g_2$, KB] = SUBSUMES* [$g_1$, $g_2$, KB*].

We then have the following:

**Theorem.** Let KB and KB* be two knowledge bases identified by the same sequence of applications of NEWKB, TELL, DEFINE and NEWKB*, TELL*, DEFINE*, respectively. Then KB is functionally equivalent to KB*.

The proof is given in an extended version of this paper (Saffiotti & Sebastiani 1988a). We notice that, as the functional-level operations are the only way a user can access a KB, the theorem actually states that the two versions of KRYPTON are indistinguishable to the user. Moreover, because we may see these operations as actually defining KRYPTON, functional equivalence is the notion of equivalence that really matters.

## 6 Extending KRYPTON

We have seen that recasting KRYPTON in terms of possible worlds semantics has the significant advantages of 1) providing (what we believe to be) a more simple and orthodox semantic account, and 2) shedding some light on the nature of the relationship between terminology and assertoricity. Nevertheless, the main driving force for this endeavour has been the extension of the language to allow the expression of modal operators, in particular those modelling terminologic belief and belief about the beliefs of multiple agents[8].

Before diving into terminologic belief and multiple agents we discuss an extension to the base language of KRYPTON that allows the expression of what we call "terminologic wffs", yielding a truly unified hybrid language. This result, beside being appealing in itself, will allow us to treat in a uniform fashion belief in propositions with either a terminologic, or an assertoric import, or both.

To obtain this we first extend the language to allow wffs of the kind (**IS** $g_1$ $g_2$), where $g_1$ and $g_2$ are gterms and **IS** is a modal operator which is intended to capture the notion of subsumption between gterms. Accordingly, we add to definition 1 the clause

$$\langle M, s \rangle \models (\mathbf{IS}\ g_1\ g_2) \text{ iff for all } t \text{ such that } s\mathcal{T}t,\ E_{\langle M, t\rangle}(g_1) \subseteq E_{\langle M, t\rangle}(g_2)$$

where $g_1$ and $g_2$ are gterms. Besides, we take (**SAME** $g_1$ $g_2$) as a metalinguistic abbreviation of (AND (**IS** $g_1$ $g_2$) (**IS** $g_2$ $g_1$)). Henceforth, we will refer to formulae of type (**IS** $g_1$ $g_2$) or (**SAME** $g_1$ $g_2$) as *terminologic wffs*.

**IS** may actually be seen as enforcing a restricted form (that is, defined on universally quantified conditionals only) of a standard "$\mu$" ("it is necessary that") modal operator[9], or, equivalently, as enforcing a restricted form (that is, defined on universally quantified formulae only) of Lewis' "strict implication". As a matter of fact, if "$\mu$" were in the language, the formula $\mu$(FORALL x (IMPLIES (P x) (Q x))) would be equivalent to (**IS** P Q).

We notice that while KRYPTON actually consists of two neatly separated languages (the language of assertoric formulae and the language of gterms), the availability of terminologic wffs renders M-KRYPTON a unified language, where assertoric and terminologic wffs have the same status. To put it another way, while in KRYPTON subsumption is a metalinguistic notion, in M-

8 This should not be taken as the statement that the original semantic account of KRYPTON given by Brachman and colleagues does not allow such an extension; we are only saying that we view PWS as a more natural framework to accomplish this.

9 This is actually a sort of return to the origins of modal logic: Aristotle, historically the first proponent of modal logic, conceived the necessity operator as qualifying not the whole sentence but the *copula* (see (Bochenski 1963:55-63) for more on Aristotle's modal logic).

KRYPTON it is a linguistic one: this means that a terminologic knowledge base may now be enforced as a logical theory in the accepted sense of the term. As a consequence, we are able to express complex formulae, such as

(OR (**IS** $g_1$ $g_2$) (**IS** $g_1$ $g_3$))

where connectives of the assertoric language apply to terminologic wffs, or

(IMPLIES (NOT (EXISTS x (AND (P x) (NOT (Q x))))) (**IS** P Q))

where both terminologic and assertoric wffs are involved. In addition, having a unified language means that it will be sufficient for the interaction language to comprise the NEWKB, ASK and TELL operators only. Notice that ASK and TELL apply now to terminologic and assertoric wffs alike: thus the effect of the KRYPTON formula DEFINE [t, g, KB] can be achieved in M-KRYPTON by TELL [(**SAME** t g), KB]; analogously, SUBSUMES [$g_2$, $g_1$, KB] becomes ASK [(**IS** $g_1$ $g_2$), KB].

Having discussed our modifications to the base language, we may now introduce belief operators and multiple believers. We are interested in the general case where there are n agents (which we call 1, ..., n), each believing in a set of propositions either of an assertoric or of a terminologic nature.

As already anticipated in previous sections, a significant body of work on formal treatments of belief has been carried out in the framework of possible worlds semantics: after the seminal works of Hintikka (1962, 1969), Kripke structures have been widely used as a formal model for a PWS-based analysis of the notion of belief (see (Halpern & Moses 1985) for a review). The intuitive idea is to associate to each agent i an accessibility relation $\mathcal{B}_i$, that can be thought as saying which worlds agent i considers as "doxastic alternatives" to a given one; that is, each of these worlds might be, *for all agent i believes*, the actual world. Hence, the formal machinery of PWS is unaltered, with the only difference being in the "semantics" of the accessibility relation: while for modal logic it formalized the notion of "ontologically" relative possibility, for doxastic logic it formalizes the notion of "doxastically" relative possibility.

Formally, we introduce into the M-KRYPTON language sentence-forming operators $\mathbf{B}_1$, ..., $\mathbf{B}_n$: if $\alpha$ is a wff then $\mathbf{B}_i\alpha$ is also a wff, to be read "agent i believes that $\alpha$". Hence, we may express "nested" beliefs, such as in $\mathbf{B}_i\mathbf{B}_j\alpha$ (when i=j we speak of "introspective belief"). In other words, an M-KRYPTON knowledge base may contain formulae representing the system's beliefs both about the domain of discourse and about the (possibly nested) beliefs of agents 1, ... , n.

To characterize our notion of belief in a formal way, we define a *hybrid Kripke structure* as a tuple M=<S, D, V, $\mathcal{T}$, $\mathcal{B}_1$, ..., $\mathcal{B}_n$>, where S, D, V and $\mathcal{T}$ are as defined in section 5 and $\mathcal{B}_1$, ..., $\mathcal{B}_n$ are binary relations over S that are serial, transitive and Euclidean[10].

Our definition of the truth relation $\models$ is then extended by means of the clause

10 We recall that a binary relation $\mathcal{B}_i$ is *serial* if, for all s∈S, there is some t such that $s\mathcal{B}_it$; is *transitive* if, for all s, t, u∈S, if $s\mathcal{B}_it$ and $t\mathcal{B}_iu$, then $s\mathcal{B}_iu$; is *Euclidean* if, for all s, t, u∈S, if $s\mathcal{B}_it$ and $t\mathcal{B}_iu$, then $t\mathcal{B}_iu$.

$\langle M, s\rangle \models (\mathbf{B}_i\ \alpha)$ iff for all t such that $s\mathcal{B}_i t$, $\langle M, t\rangle \models \alpha$

As we had anticipated, the tight integration between the terminologic and the assertoric languages that results from the introduction of the **IS** operator allows us to have a unique operator for expressing an agent's beliefs both on what the world and on what the conceptual vocabulary are like.

For better convenience, we show the definition of $\models$ in its final form:

| | | |
|---|---|---|
| $\langle M, s\rangle \models (p\ x_1 \ldots x_n)$ | iff | $\langle V(x_1), \ldots, V(x_n)\rangle \in E_{\langle M, s\rangle}(p)$ |
| $\langle M, s\rangle \models (\text{NOT}\ \alpha)$ | iff | it is not the case that $\langle M, s\rangle \models \alpha$ |
| $\langle M, s\rangle \models (\text{OR}\ \alpha\ \beta)$ | iff | $\langle M, s\rangle \models \alpha$ or $\langle M, s\rangle \models \beta$ |
| $\langle M, s\rangle \models (\text{EXISTS}\ x\ \alpha)$ | iff | $\langle M_{[x/d]}, s\rangle \models \alpha$ for some $d \in D$ |
| $\langle M, s\rangle \models (\mathbf{IS}\ g_1\ g_2)$ | iff | for all t such that $s\mathcal{T}s'$ $E_{\langle M, t\rangle}(g_1) \subseteq E_{\langle M, t\rangle}(g_2)$ |
| $\langle M, s\rangle \models (\mathbf{B}_i\ \alpha)$ | iff | for all t such that $s\mathcal{B}_i s'$, $\langle M, t\rangle \models \alpha$ |

As a consequence of these extensions, there are a number of interesting properties that hold of M-KRYPTON (besides those that already hold of KRYPTON); hereafter, we comment on the more relevant ones:

1) there is an enriched variety of formulae that are valid according to our semantics (i.e. those formulae $\alpha$ such that ASK [$\alpha$, KB] = yes for any KB): in particular, we now have valid formulae with a terminologic component such as (**IS** (ConGeneric p q) p) or (IMPLIES (**IS** p q) (FORALL x (IMPLIES (p x) (q x)))) and valid formulae involving the belief operators;
2) agents cannot hold inconsistent beliefs (this is due to the seriality of the accessibility relation) : that is, ASK* [(IMPLIES ($\mathbf{B}_i$ $\alpha$) (NOT($\mathbf{B}_i$(NOT $\alpha$)))), KB] = yes for all $\alpha$ and for all KBs;
3) agents are "conscious" of what they believe (this is due to the transitivity of the accessibility relation) : that is, ASK* [(IMPLIES($\mathbf{B}_i$ $\alpha$) ($\mathbf{B}_i$($\mathbf{B}_i$ $\alpha$))), KB] = yes for all $\alpha$ and for all KBs;
4) agents are "conscious" of what they do not believe (this is due to the Euclideanness of the accessibility relation): that is, ASK* [(IMPLIES(NOT($\mathbf{B}_i$ $\alpha$))($\mathbf{B}_i$(NOT($\mathbf{B}_i$ $\alpha$)))), KB] = yes for all $\alpha$ and for all KBs;
5) agents believe in the "assertoric counterparts" of their terminologic beliefs; for example, ASK* [($\mathbf{B}_i$(**IS** P Q), KB] = yes implies that ASK* [($\mathbf{B}_i$(FORALL x (IMPLIES (P x) (Q x))), KB] = yes for all P, Q and for all KBs;
6) agents believe in all valid propositions: that is, ASK* [($\mathbf{B}_i$(**IS** (ConGeneric P Q) P)), KB] = yes for all P, Q and for all KBs;
7) beliefs are closed wrt implication: that is, ASK* [($\mathbf{B}_i$ $\alpha$), KB] = yes and ASK* [($\mathbf{B}_i$(IMPLIES $\alpha$ $\beta$)), KB] = yes implies that ASK* [($\mathbf{B}_i$ $\beta$), KB] = yes for all $\alpha$, $\beta$ and for all KBs;
8) agents do not necessarily believe in all true propositions: that is, ASK* [$\alpha$, KB] = yes does not imply in general that ASK* [($\mathbf{B}_i$ $\alpha$), KB] = yes
9) agents may believe in false propositions: that is, ASK* [($\mathbf{B}_i$ $\alpha$), KB] = yes does not imply that ASK* [$\alpha$, KB] = yes;

10) $\mathbf{B_i}$ realizes a *de dicto* reading of belief: that is, ASK* [($\mathbf{B_i}$ $\alpha$), KB] = yes and ASK* [$\mathbf{B_i}$(**SAME** p q), KB] = yes imply that ASK* [($\mathbf{B_i}$ $\alpha_{[p/q]}$), KB] = yes even if ASK* [(**SAME** p q), KB] = no

Properties 2), 3) and 4) indicate that the $\mathbf{B_i}$ operators realize a (hybrid) version of the *KD45* doxastic logic (also known as *weak S5*). Again, from property 5) we may notice that the **IS** operator realizes a version of Feys' (1937) *T* modal logic, a version where the necessity operator is defined on universally quantified conditionals only. Properties 6) and 7) are interpreted as saying that our agents are perfect reasoners on their belief sets (that is, they are logically omniscient in the M-KRYPTON logic). Property 9) indicates that the notion modelled by the $\mathbf{B_i}$ operators is belief, not knowledge (knowledge being usually defined as justified true belief). Finally, properties 8), 9) and 10) show that the belief set of an agent is independent from what is true in the real world.

## 7 Conclusion

Our goal to investigate a notion of "awareness" that be significant both for accounting for certain situations where the "attitude bases" of agents are not closed with respect to logical consequence, and for limiting the computational load of reasoning about attitudes, has lead us to develop a language where beliefs in propositions of a terminologic nature can be expressed, and where an agent's beliefs in propositions (with either terminologic or assertoric nature) are interpreted relative to the agent's terminologic beliefs, i.e. *de dicto*. The latter is nothing but the practical realization of Bonomi's remark (see footnote 4). The belief set of an M-KRYPTON agent may thus comprise:

1) beliefs about the nature of the entities of the domain of discourse and of the conceptual relationships between them;
2) beliefs about facts about these entities;
3) beliefs about the belief sets of other agents.

This has been achieved by combining insights from the disciplines of knowledge representation and doxastic logic. In particular, we have recast KRYPTON in terms of possible worlds semantics, obtaining a semantic account that, besides being "functionally" equivalent to the original one, has allowed an easy extension to the representation of beliefs about the beliefs of multiple agents and of beliefs in propositions of a terminologic nature. Other concepts that have traditionally been the subject of modal and related logics, such as other mental attitude could be embedded in this framework.

Finally, one last word about awareness (we had left it in section 3 and never taken it up again).

First, we point out that M-KRYPTON is not a solution in itself to the problem of awareness, but was meant to be a tool for modelling it; that is, awareness was meant to be an M-KRYPTON theory (in the logical sense).

Second, a further bit of the continuing story about awareness. It went in a different, unforeseen way: M-KRYPTON actually revealed an inadequate tool, not because of wrong choices in the ways of integrating belief, terminology and assertoricity, but because of the way the terminologic language of KRYPTON had been conceived. A different tool had to be designed, that borrowed ideas from the OMEGA knowledge representation language (Attardi & Simi 1981, 1986); however, the way it integrated belief, terminology and assertoricity remained the same as in M-KRYPTON.

## Acknowledgements

This paper refers to a lecture delivered at the summer school "Programming '88" held in 1988 in Varna, Bulgaria, and includes excerpts from (Saffiotti & Sebastiani 1988b). The second author wishes to thank Peter Barnev for the invitation to take part in the school, and all participants in the school for making his stay so exciting. The authors are especially grateful to Francesco Orilia for many enlightening discussions. Work by the first author has been carried out in the frame of ESPRIT project 527 "Communication Failure in Dialogue: Techniques for Detection and Repair".

## References

Attardi, Giuseppe & Simi, Maria (1981). Semantics of inheritance and attributions in the description system Omega (Revised version). Memo n. 642, MIT AI Lab, Cambridge MA.

Attardi, Giuseppe & Simi, Maria (1986). A description oriented logic for building knowledge bases. *IEEE Proceedings 74*, pp. 1335-1344.

Barwise, Jon & Perry, John (1983). *Situations and attitudes.* Cambridge, MA: MIT Press.

Bochenski, I. Maria (1963). *Ancient formal logic.* Amsterdam NL: North Holland.

Bonomi, Andrea (1983). *Eventi mentali.* Milano I: il Saggiatore.

Brachman, Ronald J.; Fikes, Richard E. & Levesque, Hector J. (1983a). Krypton: integrating terminology and assertion. In *Proceedings of AAAI-83*, Washington DC, pp. 31-35.

Brachman, Ronald J.; Fikes, Richard E. & Levesque, Hector J. (1983b). Krypton: a functional approach to knowledge representation. *IEEE Computer 16*, pp. 67-73.

Brachman, Ronald J. & Levesque, Hector J. (eds.) (1985). *Readings in knowledge representation.* Los Altos CA: Morgan Kaufmann.

Brachman, Ronald J.; Pigman Gilbert, Victoria & Levesque, Hector J. (1985). An essential hybrid reasoning system: knowledge and symbol level accounts of KRYPTON. In *Proceedings of IJCAI-85*, Los Angeles CA, pp. 532-539.

Chellas, Brian F. (1980). *Modal logic: an introduction.* Cambridge GB: Cambridge University Press.

Cohen, Philip R. & Levesque, Hector J. (1985). Speech acts and rationality. In *Proceedings of the 23rd Annual Meeting of the ACL,* Chicago IL, pp. 49-60.

Fagin, Ronald & Halpern, Joseph Y. (1988). Belief, awareness and limited reasoning. *Artificial Intelligence 34*, pp. 39-76.

Feys, Robert (1937). Les logiques nouvelles des modalités. *Revue Neoscholastique de Philosophie 40*, pp. 517-553.

Halpern, Joseph Y. & Moses, Yoram O. (1984). Knowledge and common knowledge in a distributed environment. In *Proceedings of the 3rd ACM Conference on Principles of Distributed Computing*, Vancouver, Canada, pp. 50-61.

Halpern, Joseph Y. & Moses, Yoram O. (1985). A guide to the modal logics of knowledge and belief: preliminary draft. In *Proceedings of IJCAI-85*, Los Angeles CA, pp. 480-490.

Hintikka, K. Jaakko (1962). *Knowledge and belief. An introduction to the logic of the two notions.* Ithaca NY: Cornell University Press.

Hintikka, K. Jaakko (1969). Semantics for propositional attitudes. In Linsky, Leonard (ed.), *Reference and modality*, London GB: Oxford University Press, 1971, pp. 145-167.

Hintikka, K. Jaakko (1975). Impossible possible worlds vindicated. *Journal of Philosophical Logic 4*, pp. 475-484.

Hirst, Graeme (1981). *Anaphora in Natural Language Understanding: A Survey.* New York NY: Springer.

Hughes, George E. & Cresswell, Maxwell J. (1968). *An introduction to modal logic.* London GB: Methuen.

Kripke, Saul A. (1963). Semantical analysis of modal logic. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik 9*, pp. 67-96.

Levesque, Hector J. (1984a). A logic of implicit and explicit belief. In *Proceedings of AAAI-84*, Austin TX, pp. 198-202 (revised version appears as Fairchild Technical Report 653, Palo Alto CA).

Levesque, Hector J. (1984b). Foundations of a functional approach to knowledge representation. *Artificial Intelligence 23*, pp. 155-212.

Perrault, C. Raymond (1987). An application of default logic to speech act theory. Technical report CSLI-87-90, Center for the study of language and information, Stanford CA.

Pollack, Martha E. (1986). A model of plan inference that distinguishes between the beliefs of actors and observers. In *Proceedings of the 24th Meeting of the ACL*, New York NY, pp. 207-214.

Reilly, Ronan (ed.) (1987). *Communication failure in dialogue and discourse: detection and repair processes.* Amsterdam NL: North Holland.

Saffiotti, Alessandro & Sebastiani, Fabrizio (1988). Dialogue modelling in M-KRYPTON, a hybrid language for multiple believers. In *Proceedings of the 4th IEEE Conference on Artificial Intelligence Applications*, San Diego, CA, pp. 56-61.

Saffiotti, Alessandro & Sebastiani, Fabrizio (1989). Towards a hybrid logic of acquaintance and awareness. *Proceedings of AISB'89, Seventh Conference of the Society for the Study of Artificial Intelligence and Simulation of Behaviour*, University of Sussex, pp. 95-103.

Searle, John R. (1975). Indirect speech acts. In Cole, Peter & Morgan, Jerry L. (eds.), *Syntax and semantics - Speech acts*, New York NY: Academic Press, pp. 59-82.

Vilain, Marc (1985). The restricted language architecture of a hybrid representation system. In *Proceedings of IJCAI-85*, Los Angeles CA, pp. 548-551.

Wahlster, Wolfgang & Kobsa, Alfred (1986). Dialogue-based user models. *IEEE Proceedings 74*, pp. 948-960.