A Hybrid Framework for Representing Uncertain Knowledge

Alessandro Saffiotti*

IRIDIA - Université Libre de Bruxelles
Av. F. Roosvelt 50 - CP 194/6
B-1050 Bruxelles - Belgium
E-mail: r01507@bbrbfu01.bitnet

Abstract

This paper addresses the problem of bridging the gap between the fields of Knowledge Representation (KR) and Uncertain Reasoning (UR). The proposed solution consists of a framework for representing uncertain knowledge in which two components, one dealing with (categorical) knowledge and one dealing with uncertainty about this knowledge, are singled out. In this sense, the framework is "hybrid". This framework is characterized in both model-theoretic and proof-theoretic terms. State of belief is represented by "belief sets", defined in terms of the "functional approach to knowledge representation" suggested by Levesque. Examples are given, using first order logic and (a minimal subset of) M-Krypton for the KR side, and a yes/no trivial case and Dempster-Shafer theory for the UR side.

1. Introduction

An impressive work has been carried out over the last two decades in the fields of Knowledge Representation (KR) and of Uncertain Reasoning (UR), resulting in a number of concepts being investigated, a number of problems being identified, and a number of solutions being developed (see e.g. Israel & Brachman, 1981; Brachman & Levesque, 1985; Saffiotti, 1987; Henkind & Harrison, 1988). Yet, curiously enough, these two fields, which we would expect to be tightly related, apparently suffered from a lack of cross-fertilization, or even of communication, during their development. As a matter of fact, the literature in each field seems to have scarcely taken into account the problems and the results emerged in the other one. As an example, consider the assertions "Birds are animals", "Typically birds fly", "Most of my friends like music" and "Smoke suggests fire": a person working in KR would recognize them as pertaining to qualitatively different types of knowledge, and would claim for different mechanisms to represent (part of) them. Though, a Bayesian would probably code all of them by the single structure $P(A|B) = x$.

In this paper we present an attempt at bridging the gap between the UR and the KR fields, by describing a formal unifying framework for representing uncertain knowledge. It is of basic importance here to make clear what we mean by "uncertain knowledge". We postulate that uncertain knowledge is composed of categorical knowledge accompanied by information regarding the uncertainty about it; for instance, the fragment of uncertain knowledge expressed by "Smoke suggests fire" can be seen as a piece of categorical knowledge (e.g. a "symptom-of" link between the concepts "Smoke" and "Fire" in a semantic net), plus information about its uncertainty (e.g. a probability value for that link). From an AI viewpoint, this corresponds to seeing uncertainty as metaknowledge. Throughout this paper, we will use the expressions "knowledge component" and "uncertainty component" to refer to the two components of uncertain knowledge. Two hypotheses are hidden in this notion of uncertain knowledge. First, knowledge, and the reasoning processes based on it, is categorical: it is the *validity* of knowledge (and that of the conclusions drawn from it) to be a matter of degree, not the knowledge itself. Second, the only uncertainty we talk about is epistemic uncertainty: uncertainty about the validity of our knowledge with respect to an intrinsically certain reality. An interesting problem is whether other types of uncertainty (vagueness, for instance) are captured by our notion or not. For the case of vagueness, the answer is affirmative, if we accept the interpretation of the vagueness of the sentence "Enzo is rich" as pertaining to the adequacy (in our mind) of the description "rich" to the individual "Enzo", rather than to the ontological fuzziness of the predicate "rich" (cf. Schefe, 1980).

The formal framework that we define in this paper deals with uncertain knowledge by singling out the knowledge component and the uncertainty component. Both the differences and the relationships between these two components are accounted for. Under the hypothesis of considering uncertainty as a kind of knowledge (about our knowledge), the proposed framework fits the "hybrid knowledge representation" paradigm (Brachman & Levesque, 1982). Apart from its theoretical interest, this framework is meant to form a basis for defining —given a KR system and a UR system— a combined uncertain

* This research has been partially supported by the ARCHON project, funded by grants from the Commission of the European Communities under the ESPRIT-II Program, P-2256. The partners in the ARCHON project are: Krupp Atlas Elektronik, Amber, CERN, CNRG-NTUA Athens, ECRC, Framentec, FWI Univ. of Amsterdam, Iberduero, Univ. Libre de Bruxelles, JRC Ispra, Labein, Volmac, Univ. of Porto, Queen Mary College.

knowledge representation system able to perform uncertain reasoning on structured knowledge.

The rest of this paper is organized as follows: section 2 presents the conceptual background and the formal definition of our framework. Section 3 gives a model-theoretic account of this framework in terms of possible worlds, and shows three examples. Section 4 mimics section 3, but uses a proof-theoretic perspective. Some hints for building a real hybrid system for uncertain knowledge are also given. Finally, section 5 concludes.

2. Hybrid Belief Structures

Once we have clear the notion of "uncertain knowledge" that we want to model, we can state two essential requirements for a general framework in which this uncertain knowledge can be represented and dealt with:

1. it must account for both the distinction and the relationship between "knowledge" and "uncertainty", as the two basic components of uncertain knowledge;
2. it must be general enough to accommodate a number of KR languages and of UR calculi: the framework should not make strong hypotheses on the form in which knowledge and uncertainty are represented.

In order to guarantee ourselves a general formalization, we adhere to what Levesque (1984) has called "functional approach to Knowledge Representation". In this approach, knowledge is represented by abstract data structures ("Knowledge Bases") characterized in terms of the operations that can be made on them. Typical operations will be a query operation "Ask", an updating operation "Tell", and an initialization operation "Empty". Correspondingly, we model uncertain knowledge in terms of abstract "Belief Sets", the uncertain correspondent of Levesque's Knowledge Bases[1]. In them, knowledge and uncertainty about this knowledge are represented according to a specific KR language and a specific UR calculus, respectively. We will write $\mathcal{BS}$ to refer to the set of belief sets. Belief sets will be characterized in terms of the following operations:

$$\text{Ask}: \mathcal{L} \times \mathcal{BS} \rightarrow \mathcal{T}$$
$$\text{Tell}: \mathcal{L} \times \mathcal{T} \times \mathcal{BS} \rightarrow \mathcal{BS}$$
$$\text{Empty}: \{\emptyset\} \rightarrow \mathcal{BS}$$

where $\mathcal{L}$ is the language used to represent knowledge (a KR language), and $\mathcal{T}$ is the set of *belief judgements* used to represent uncertainty about this knowledge (numbers, "true/false", tokens, etc.). Intuitively, "Ask[α,κ]" returns an element of $\mathcal{T}$ representing the extent to which the belief expressed by formula α of $\mathcal{L}$ may be considered true[2] in the belief set κ, while "Tell[α,x,κ]" returns the new belief set obtained by assimilating in κ the belief expressed by α with belief judgement x. "Empty[]" simply returns a knowledge base without any knowledge.

Unfortunately, the operations proposed by Levesque do not account for the dichotomous nature of uncertain knowledge. We want to find a decomposition of these operations, in which the KR component and the UR component are singled out. In order to do this, we will borrow from another well-known dichotomy: the distinction between extension and intension. By *intension* of a sentence we mean its "meaning", i.e. the idea expressed by it. By *extension* of a sentence we mean the particular object designated by it (i.e. a truth value). E.g. the two sentences "17 is a prime" and "A=A" have the same extension (*true*), but different intensions. For our goals, we accept the following characterization of intensions: the intension of a sentence is a function from possible states of affair (contexts) to the truth value of the sentence in each context. In this light, we can state the problem addressed by the Ask operation as that of judging how much the context (partially and doubtfully) specified by a belief set κ is likely to be one in which the asked formula is true. We decompose this problem into two steps

1. Given formula α, evaluate its intension (written $\|\alpha\|$);
2. Given $\|\alpha\|$, evaluate its belief judgement in the context (partially and doubtfully) specified by κ.

Intuitively, step 1 corresponds to a process of meaning attribution, i.e. a mapping between the structures we use to represent knowledge and the knowledge which is represented. This process is typically performed by a KR system. So step 1 identifies the KR component in the Ask process. On the other hand, in step 2 we completely disregard linguistic structures: rather, we evaluate belief judgements of abstract propositions (intensions) in a given context. This corresponds to the entailment part of what has been traditionally called "the evidential reasoning problem" in the UR literature (e.g. Thompson, 1985). Thus, step 2 identifies the UR component of the Ask process. An analogous decomposition of "Tell(α,x,κ)" can be given as: 1) evaluate $\|\alpha\|$ (meaning attribution), and 2) update κ by assimilating $\|\alpha\|$ (with its belief judgement x) in it (updating part of the evidential reasoning problem). The above decompositions are graphically summarized below

$$\begin{tikzcd}[row sep=small] \text{Ask}: & \mathcal{L} \arrow[d, "\Phi"'] & \times & \mathcal{BS} \arrow[d, "I"'] & & & \text{Tell}: & \mathcal{L} \arrow[d, "\Phi"'] & \times & \mathcal{T} \arrow[d, "I"'] & \times & \mathcal{BS} \arrow[d, "I"'] & \\ & \mathcal{I} & \times & \mathcal{BS} \arrow[r, "\Psi"] & \mathcal{T} & & & \mathcal{I} & \times & \mathcal{T} & \times & \mathcal{BS} \arrow[r, "\Psi"] & \mathcal{BS} \end{tikzcd}$$

where $\mathcal{I}$ is the set of intensions, Φ is a function responsible for attributing a meaning to formulae, Ψ is a function responsible for solving the (two facets[3] of the) evidential reasoning problem, and *I* is the identity.

[1] Here "belief", as opposed to "knowledge", is not related to what is true in the real world, and is a matter of degree. We later use "belief judgement" instead of "truth value" in a similar way.

[2] Notice that we are only considering yes-no queries; this assumption is not restrictive (Levesque, 1984).

[3] Ψ will be actually considered as a pair of functions.

We are now in a position to formally define hybrid belief structures, the basic ingredient of our framework. As a notational convention, we use α, β,... to denote formulae of $\mathfrak{L}$; x, y,... for belief judgements; p, q,... for intensions; κ_1, κ_2,... for belief sets; and $\wp\,\mathfrak{A}$ for the power set of set $\mathfrak{A}$.

Def. 1. *Let $\mathfrak{L}$ be a language, and $\mathfrak{T}$ and $\mathfrak{I}$ be non empty sets. Let $\mathfrak{B}$ stand for $(\mathfrak{I} \times \mathfrak{T})$. Let also Φ be a function from $\mathfrak{L}$ to $\mathfrak{I}$, and Ψ be a pair of functions (Ψ^1,Ψ^2) such that $\Psi^1\colon \mathfrak{I} \times \wp\,\mathfrak{B} \rightarrow \mathfrak{T}$ and $\Psi^2\colon \mathfrak{B} \times \wp\,\mathfrak{B} \rightarrow \wp\,\mathfrak{B}$. Then we call the tuple $\mathfrak{H} = \langle \mathfrak{L}, \Phi, \mathfrak{I}, \mathfrak{T}, \Psi \rangle$ a hybrid belief structure.*

In words, we are given a KR language $\mathfrak{L}$, and a function Φ for it which maps each formula of $\mathfrak{L}$ to an intension (its meaning). $\mathfrak{L}$ and Φ together constitute the *KR component* of $\mathfrak{H}$. From the other hand, we are given the set $\mathfrak{T}$ of belief judgements of an UR calculus, and a pair of functions (Ψ^1,Ψ^2) which define its dynamic behaviour on the space of the intensions. Ψ^1 returns a belief judgement given an intension and a belief set: this is the belief judgement for the given intension with respect to the belief judgements of the other beliefs in the belief set. Ψ^2 returns a (new) belief set given a (new belief judgement for a) belief and a (old) belief set: this is the updated belief set in which the belief judgements of all beliefs have been modified in order to account for the new one. $\mathfrak{T}$ and Ψ together constitute the *UR component* of $\mathfrak{H}$. Notice that the only hypothesis made on the KR and the UR components is that their semantics can be described in terms of the functions Φ and Ψ above. These components act as independent and autonomous systems, which cooperate via elements of $\mathfrak{I}$. Intensions are used as abstract objects which represent our knowledge independently from the particular linguistic structures we use (in $\mathfrak{L}$) to express it: it is to these objects that uncertainty is allocated.

We now define belief sets trough the operations which characterize their behaviour:

Def. 2. *Let $\mathfrak{H} = \langle \mathfrak{L},\Phi,\mathfrak{I},\mathfrak{T},\Psi \rangle$ be a hybrid belief structure. Then belief sets on $\mathfrak{H}$ are defined by the following operations:*

$$\begin{aligned}
\mathbf{Empty}_{\mathfrak{H}}[\,] &= \{ \langle \Phi(\underline{\mathrm{true}_{\mathfrak{L}}}), \underline{\mathrm{true}_{\mathfrak{T}}} \rangle \} \\
\mathbf{Ask}_{\mathfrak{H}}[\alpha, \kappa] &= \Psi^1(\Phi(\alpha), \kappa) \\
\mathbf{Tell}_{\mathfrak{H}}[\alpha, x, \kappa] &= \Psi^2(\langle \Phi(\alpha), x \rangle, \kappa)
\end{aligned}$$

where $\underline{\mathrm{true}_{\mathfrak{L}}}$ represents the tautology of $\mathfrak{L}$, and $\underline{\mathrm{true}_{\mathfrak{T}}} \in \mathfrak{T}$ represents total confidence.

Operationally speaking, belief sets are built starting from $\mathrm{Empty}_{\mathfrak{H}}[\,]$, and then by performing successive $\mathrm{Tell}_{\mathfrak{H}}$ operations on it. Thus, belief sets are sets of pairs $\langle p,x \rangle$, where p is an intension and $x \in \mathfrak{T}$. Notice that, while we use formulae of $\mathfrak{L}$ to interact with belief sets via the Ask and Tell operations, belief judgements are actually associated with the intensions connoted by these formulae (and not to the formulae themselves).

3. The Semantic Perspective

We now want to see how the belief sets defined above can actually perform uncertain reasoning on structured knowledge, possibly merging together already existing KR and UR systems. We will make a two step descent from the abstract level we were before: as a first step, we will consider particular choices for the set of intensions $\mathfrak{I}$. This will of course pose more constraints on the form of the Φ and Ψ functions, and then on the set of KR and UR calculi we can capture. As a second step, we will consider full instantiations, where the elements of the framework will be completely specified for a particular choice of KR and UR calculi. In this section, we stick to a model-theoretic approach to both KR and UR, by using possible worlds to represent intensions. In the next section, we will use (sets of sets of) formulae to represent intensions, so switching to a syntactic viewpoint.

Let $M = \langle S, D, V, \{\mathfrak{R}_i \mid i \geq 0\} \rangle$ be a Kripke structure[4], where as usual S is a set of states, D is a domain of individuals, V is a mapping from symbols of $\mathfrak{L}$ and states $s \in S$ to elements (and sets) of D, and the $\mathfrak{R}_i$'s are binary relations over S. A Kripke world is a pair $\langle M,s \rangle$, with $s \in S$. We then let

$$\mathfrak{I} \subseteq \wp\{\langle M,s \rangle \mid \langle M,s \rangle \text{ is a Kripke world}\}$$

Given a KR language $\mathfrak{L}$, and a formula α of it, $\Phi(\alpha)$ consists of the set of worlds where α holds according to the semantics given to $\mathfrak{L}$. Thus, according to this semantic perspective, we must be given a model-theoretic account of any KR system we want to fit in a hybrid belief structure, where the notion of truth in a world is defined. A belief is a set of worlds together with a belief judgement, and belief sets are composed of these beliefs. The pair of functions Ψ must consider sets of worlds as its basic objects. This means that a description of an UR technique in terms of possible worlds must be available in order to use this technique in a hybrid belief structure.

Example 1: First Order Logic + {yes, no}
Only for the sake of familiarizing the reader with the use of hybrid belief structures and of Ask and Tell operations to define belief sets, we present a very simple example, where knowledge is represented in standard first order logic (FOL), and uncertainty is represented by yes/no values. We need to define

[4] Other mathematical structures which are used to give semantics to KR languages could have been employed. Kripke structures have been chosen here mainly because of their wide use within the AI community (e.g. Halpern & Moses, 1985).

the elements of $\mathcal{H}_1 = \langle \mathcal{L}_{FOL}, \Phi_{FOL}, \mathcal{I}, \mathcal{V}_{TF}, \Psi_{TF} \rangle$. $\mathcal{L}_{FOL}$ is a standard first order language, and $\mathcal{V}_{TF}=\{\mathit{yes}, \mathit{no}\}$. $\mathcal{I}$ is composed of sets of worlds $\langle M,s_0 \rangle$ such that $M = \langle \{s_0\},D,V \rangle$ is a standard FOL interpretation structure. We indicate by W the set of all these worlds. We then define

$$\Phi_{FOL}(\alpha) = \{ \langle M,s_0 \rangle \mid \langle M,s_0 \rangle \models_{FOL} \alpha \}$$

where $\models_{FOL}$ is the standard truth relation for FOL. In words, $\Phi_{FOL}(\alpha)$ returns the set of all the first order models of α.

The definition of the Ψ_{TF} pair of functions is given by:

$$\Psi^1_{TF}(p, \kappa) = \begin{cases} \mathit{yes} & \text{if for all } \langle q, \mathit{yes} \rangle \in \kappa,\ q \subseteq p \\ \mathit{no} & \text{otherwise} \end{cases}$$

$$\Psi^2_{TF}(\langle p,x \rangle, \kappa) = \begin{cases} \{\langle q \cap p, \mathit{yes} \rangle \mid \langle q, \mathit{yes} \rangle \in \kappa\} & \text{if } x=\mathit{yes} \\ \kappa & \text{otherwise} \end{cases}$$

The condition "$q \subseteq p$" can be read in terms of logical entailment: given two formulae α and β of $\mathcal{L}$, respectively connoting p and q, $q \subseteq p$ is true whenever $\{w \mid w \models \beta\} \subseteq \{w \mid w \models \alpha\}$, that is whenever $\beta \models \alpha$. As for the updating side, the belief set obtained from κ by assigning the value **true** to the intension p is made by retaining those sets of worlds in κ which are consistent with p. Belief set operations are then easily defined in terms of Φ_{FOL} and Ψ_{TF} [5]:

$$\mathbf{Empty_1}[] = \{\langle W, \mathit{yes} \rangle\}$$

$$\mathbf{Ask_1}[\alpha, \kappa] = \begin{cases} \mathit{yes} & \text{if for all } \langle q, \mathit{yes} \rangle \in \kappa,\ q \subseteq \Phi_{FOL}(\alpha) \\ \mathit{no} & \text{otherwise} \end{cases}$$

$$\mathbf{Tell_1}[\alpha, x, \kappa] = \begin{cases} \{\langle q \cap \Phi_{FOL}(\alpha), \mathit{yes} \rangle \mid \langle q, \mathit{yes} \rangle \in \kappa\} & \text{if } x=\mathit{yes} \\ \kappa & \text{otherwise} \end{cases}$$

As required, knowledge is expressed in these belief sets through FOL formulae, and uncertainty about it is expressed by "yes/no" values. For instance, if the belief set κ is built by

$\mathrm{Tell_1}[\mathrm{dog(Alex)}, \mathit{yes}, \mathrm{Tell_1}[\forall x.(\mathrm{dog}(x) \supset \mathrm{animal}(x)), \mathit{yes}, \mathrm{Empty_1}[]]]$

then we have

$\mathrm{Ask_1}[\mathrm{animal(Alex)}, \kappa] = \mathit{yes}$

Notice that the definition of Ask accounts for the incompleteness of belief sets: e.g. $\mathrm{Ask_1}[\alpha, \mathrm{Empty_1}[]] = \mathit{no}$ and $\mathrm{Ask_1}[\sim\alpha, \mathrm{Empty_1}[]] = \mathit{no}$ for every non valid formula α.

Example 2: FOL + Dempster-Shafer Theory

In the next example we want to define belief sets in which knowledge is represented in FOL, while uncertainty is dealt with according to the Dempster-Shafer (DS) theory of evidence (Shafer, 1976; Smets, 1988). As you could expect, only the UR component of the $\mathcal{H}_1$ defined in the previous example needs to be modified. Our new hybrid belief structure is $\mathcal{H}_2 = \langle \mathcal{L}_{FOL}, \Phi_{FOL}, \mathcal{I}, \mathcal{V}_{DS}, \Psi_{DS} \rangle$, where $\mathcal{L}_{FOL}$, Φ_{FOL} and $\mathcal{I}$ are as in $\mathcal{H}_1$. $\mathcal{V}_{DS}$ is now the [0,1] interval. A belief set on $\mathcal{H}_2$, then, is a set of pairs $\langle p,x \rangle$, with p being a set of worlds and $x \in [0,1]$. In the terminology of the DS theory, we can see a belief set κ as a "basic probability assignment" (bpa) on the set W of possible (FOL) worlds, i.e. a distribution of a unitary mass among the subsets of W.

$\Psi^1_{DS}(p,\kappa)$ returns an element of [0,1] that corresponds, in DS terminology, to the value of Bel(p) with respect to the bpa represented by κ:

$$\Psi^1_{DS}(p, \kappa) = \sum_{\substack{\langle q,x \rangle \in \kappa \\ q \subseteq p}} x$$

Intuitively, our confidence in the truth of the knowledge represented by p is just the sum of the mass values attributed to any intension in κ whose truth entails the truth of p. As for the updating half of Ψ_{DS}, we notice that updating is typically performed in DS theory by combining, according to the so-called Dempster's rule of combination, the bpa representing the present state of belief with that representing the new evidence. So, we combine the old bpa κ with a bpa which allocates the desired amount of belief to the new intension[6], and the rest to W.

$$\Psi^2_{DS}(\langle p, x \rangle, \kappa) = \kappa \oplus \{\langle p, x \rangle, \langle W, 1-x \rangle\}$$

where $\oplus$ stands for Dempster's rule of combination, recast in terms of belief sets (see full paper). The definitions of $\mathrm{Empty_2}$, $\mathrm{Ask_2}$ and $\mathrm{Tell_2}$ are then given by:

$$\mathbf{Empty_2}[] = \{\langle W, 1 \rangle\}$$

$$\mathbf{Ask_2}[\alpha, \kappa] = \sum \{x \mid \langle q,x \rangle \in \kappa \text{ and } q \subseteq \Phi_{FOL}(\alpha)\}$$

$$\mathbf{Tell_2}[\alpha, x, \kappa] = \kappa \oplus \{\langle \Phi_{FOL}(\alpha), x \rangle, \langle W, 1-x \rangle\}$$

We show the behaviour of these belief bases by a simple example. Consider the set of worlds $W=\{\emptyset,d,a,da\}$, where D contains the only individual <u>alex</u>, and V is such that in the world "d", <u>alex</u> is in the extension of <u>Dog</u>, but not in that of <u>Animal</u>; in "da", he is in the extension of both <u>Dog</u> and <u>Animal</u>, etc. We then have the following belief sets:

$\kappa_0 \equiv \mathbf{Empty_2}[] = \{\langle W, 1 \rangle\}$

$\kappa_1 \equiv \mathbf{Tell_2}[(\underline{\mathrm{isa}}\ \mathrm{Dog\ Animal}), 0.9, \kappa_0] = \{\langle W, 0.1 \rangle, \langle \{\emptyset,a,da\}, 0.9 \rangle\}$

$\kappa_2 \equiv \mathbf{Tell_2}[(\mathrm{Dog\ alex}), 0.7, \kappa_1] = \{\langle W, 0.03 \rangle, \langle \{d,da\}, 0.07 \rangle, \langle \{\emptyset,a,da\}, 0.27 \rangle, \langle \{da\}, 0.63 \rangle\}$

From what we get, for instance:

$\mathbf{Ask_2}[(\mathrm{Animal\ alex}), \kappa_2] = 0.63$

Example 3: "ZERO" KR System + DS Theory

In the last example of this section we use a KR system, called ZERO, which, though almost unreasonably simple, presents some interesting characteristics. First, ZERO incorporates concepts —like the distinction between "facts" and "definitions"— that are traditionally addressed in the KR field, but which are not even taken into consideration in the UR tradition. The resulting belief sets will therefore exhibit conceptually non trivial (though minimal) KR capabilities, together with the power of the DS approach in dealing with uncertainty. Second, ZERO constitutes the very core of the M-KRYPTON KR language[7]:

---

[5] Not surprisingly, the definitions of Ask and Tell for this "collapsed" case basically corresponds to those in (Levesque, 1984).

[6] Notice that we are using in the Tell primitive a single value x to specify confidence in the truth of α. We can easily generalize x to pairs (x_t, x_f), where x_t and x_f measure our confidence in the truth of α, and in its falseness, respectively. The full formalization of Dempster-Shafer belief sets is given in the fool paper.

[7] M-Krypton (Saffiotti & Sebastiani, 1988) is a KR language that extends Krypton (Brachman et al. 1985) with belief operators for

we could use the same procedure described in this example to define belief sets with M-KRYPTON as KR component (by just replacing the truth function $\models_Z$ of ZERO by that of M-KRYPTON). Third, ZERO has been implemented and used as a test-bed to experiment with the ideas presented here.

Formulae of ZERO are of two kinds[8]: *definitions*, of form (isa C_i C_j), to be read as "C_i is a sub-concept of C_j", and *facts*, of form (C_i a_j), to be read as "a_j is an instance of concept C_i". Semantics is given to ZERO in terms of Kripke structures in the following way:

$\langle M,s\rangle \models_Z (C_i\ a_j)$ *iff* $V(a_j, s) \in V(C_i, s)$

$\langle M,s\rangle \models_Z (\underline{isa}\ C_i\ C_j)$ *iff* *forall* s' *s.t.* $s\mathcal{R}s'$, $V(C_i,s') \subseteq V(C_j,s')$

The use of the $\mathcal{R}$ relation guarantees the one-way dependence between facts and definitions: if (isa C_1 C_2) holds in a world, all instances of C_1 in that world will necessarily be instances of C_2 as well; yet, the opposite implication does not hold.

Our new hybrid belief structure is $\mathcal{H}_3 = \langle \mathcal{L}_Z, \Phi_Z, \mathcal{I}, \Upsilon_{DS}, \Psi_{DS}\rangle$, where Υ_{DS}, Ψ_{DS} and $\mathcal{I}$ are as in $\mathcal{H}_2$, and $\mathcal{L}_Z$ is the language of ZERO. The definition of Φ_Z is plainly:

$$\Phi_Z(\alpha) = \{ \langle M,s\rangle \mid \langle M,s\rangle \models_Z \alpha \}$$

$\mathbf{Empty}_3$, $\mathbf{Ask}_3$ and $\mathbf{Tell}_3$ are exactly the same as in the preceding example, with Φ_Z replacing Φ_{FOL}. As a result, the modelled belief sets will represent knowledge using the ZERO language and mechanisms, and uncertainty using DS. The following example illustrates the behaviour of these belief sets:

$\kappa_1 \equiv$ **Tell$_3$**[(isa Researcher Brontosaurus), 0.7, **Empty$_3$**[]]

$\kappa_2 \equiv$ **Tell$_3$**[(isa Brontosaurus Animal), 0.9, κ_1]

$\kappa_3 \equiv$ **Tell$_3$**[(Researcher alex), 1, κ_2]

From this we get, for instance:

Ask$_3$[(Brontosaurus alex), κ_3] = 0.9

Ask$_3$[(Animal alex), κ_3] = 0.63

4. The Syntactic Perspective

Though convenient from a formal point of view, the above description in terms of manipulation of abstract possible worlds does not shade much light on the practical side of our enterprise. The hybrid belief structure formalism, however, is not constrained to use sets of possible worlds to represent intensions. As proof-theoretic descriptions are normally more convenient than model-theoretic ones whenever concreteness is at issue, we move now to a proof-theoretic approach, and consider representing intensions in terms of more "tangible" syntactical structures, rather than semantical ones. Given a KR language $\mathcal{L}$, equipped with a proof theory, we define a *possible argument* in $\mathcal{L}$ to be a consistent set of formulae of $\mathcal{L}$. We let intensions be sets of possible arguments:

$$\mathcal{I}' \subseteq \wp(\wp\mathcal{L})$$

Given a KR language $\mathcal{L}$, and a formula α of it, we focus on all those possible arguments in $\mathcal{L}$ such that α is deducible from them according to the given proof-theory. We then let $\Phi(\alpha)$ consist of the collection of all these sets of formulae[9]. Therefore, according to this syntactic perspective, we must be given a proof-theoretic account of any KR system we want to fit in a hybrid belief structure, where the notion of deduction is defined (but notice that we do not need a truth relation for $\mathcal{L}$ any more). A belief is then a set of possible arguments together with a belief judgement, and belief sets are composed of these beliefs.

Example 4: "ZERO" KR System + DS

We restate example 3 in proof-theoretic terms. ZERO provides us with a proof theory (and hence a deduction operator $\vdash_Z$) consisting of one axiom schemata and two inference rules:

A1. (isa C_i C_i)

R1. *From* (C_i a_k) *and* (isa C_i C_j) *infer* (C_j a_k)

R2. *From* (isa C_i C_j) *and* (isa C_j C_k) *infer* (isa C_i C_k)

Our new hybrid belief structure is $\mathcal{H}_4 = \langle \mathcal{L}_Z, \Phi'_Z, \mathcal{I}', \Upsilon_{DS}, \Psi_{DS}\rangle$. Apart from $\mathcal{I}'$, the only difference with $\mathcal{H}_3$ is in the definition of Φ_Z, given now in terms of deduction:

$$\Phi'_Z(\alpha) = \{ \pi \in \boldsymbol{P} \mid \pi \vdash_Z \alpha \}$$

where $\boldsymbol{P}$ is the set of all possible possible arguments for ZERO. As for the uncertainty component of $\mathcal{H}_4$, Υ_{DS} and Ψ_{DS} are exactly the same as in $\mathcal{H}_3$. The intuitive interpretation is however different. The condition $\mathbf{q} \subseteq \mathbf{p}$ in Ψ^1_{DS} should now be read in terms of deduction: $\{\pi \mid \pi \vdash \beta\} \subseteq \{\pi \mid \pi \vdash \alpha\}$ is true when, whatever argument is valid among those proving β, it proves α as well, i.e. whenever $\beta \vdash \alpha$. The Ψ^2_{DS} function, by considering the intersection of possible arguments (via Dempster's rule), extends the set of focal elements to include the intension corresponding to the conjunction of all the knowledge already in κ with the new one.

Ask$_4$ and **Tell$_4$** are the same as in example 3, with Φ' replacing Φ, while **Empty$_4$**[] = $\{\langle \boldsymbol{P},1\rangle\}$. The behaviour of the resulting belief sets on $\mathcal{H}_4$ will be the same as in the $\mathcal{H}_3$ case: using a model-theoretic or a proof-theoretic approach is transparent to the user of the resulting belief sets.

Going down to the syntactic level gets us closer to the implementation side. Consider the architecture below:

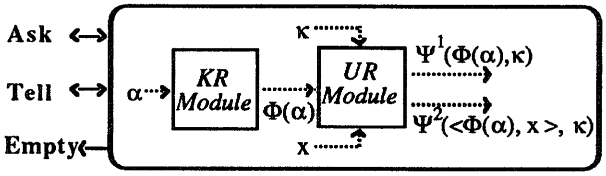

multiple agents. More important here, it provides a Kripke semantics for Krypton in the same style of that given to ZERO.

[8] A full description of ZERO is given in the full paper.

[9] When $\mathcal{L}$ comprises a conjunction operator $\wedge$ satisfying "$\{\alpha,\beta\} \vdash \gamma$ iff $\{\alpha\wedge\beta\} \vdash \gamma$", and a disjunction operator $\vee$ satisfying "$\{\alpha\} \vdash \gamma$ or $\{\beta\} \vdash \gamma$ iff $\{\alpha\vee\beta\} \vdash \gamma$", a (finite) set of possible arguments $\{\pi_i \mid \pi_i = \{\beta_{i1},\dots,\beta_{iN_i}\}\}$ is equivalent to the formula of $\mathcal{L}$ "$\vee_i(\wedge_j \beta_{ij})$". In such cases, we could represent intensions by formulae of $\mathcal{L}$. However, this is not true in general (eg Zero does not fit this case).

The role of the KR module (possibly an already existing system) is to compute $\Phi(\alpha)$ for each formula α of the KR language. When $\Phi(\alpha)$ is a set of possible arguments, this role is not so far from a plausible behaviour for a KR system. Though, this set includes all the possible arguments for α according to the deduction theory of $\mathfrak{L}$, and this is in general a computationally intractable object. Two steps may be undertaken at this stage:

1) to only consider possible arguments which are "reasonable" with respect to what is actually believed in the belief sets; and
2) to let the KR module provide "fragments" of possible arguments, generated while performing inferences; the reconstruction of full possible arguments from these fragments is then performed by the UR module.

Step 1 means that the KR module must access the content of the belief set to decide which inferences to draw. Step 2 greatly weakens the demands on the KR module. Moreover, if the fragments above correspond to reports of single inference steps, they can be seen as ATMS justifications (deKleer, 1986). We can then use an ATMS inside the UR module to reconstruct the full possible arguments (ATMS "environments") given these justifications[10]. The discussed architecture, and the two steps above, are detailed in (Saffiotti, 1990), where a possible algorithm based on an ATMS is also proposed.

5. Conclusions and Related Work

We have presented an approach to link an arbitrary model for UR to an arbitrary KR system (provided that suitable formalizations are available for both). This approach has been formalized by defining in a functional way belief sets, abstract data types where uncertain knowledge is represented as knowledge (dealt with by the KR system) plus uncertainty (dealt with by the UR calculus). Belief sets are peculiar in that they associate uncertainty to the knowledge itself rather than to the linguistic structures used to represent it. From the point of view of KR, belief sets are a tool for attaching an arbitrary treatment of uncertainty to a KR system. From the point of view of UR, they constitute a tool for extending the applicability of an UR technique to kinds of knowledge that would otherwise be difficult to express in the language of the mathematical model of the UR calculus. E.g. consider expressing in a standard DS formalism

Tell[$\forall$x.($\exists$y.child(x,y)$\rightarrow$married(x)), 0.9, **Empty**[]]

The idea of marring uncertain reasoning with knowledge representation seems to be fairly new in the literature. Some symptom of this tendency may be found in (Zadeh, 1989). However, Zadeh's solution consists in proposing Fuzzy Logic as a KR tool, while we suggest to combine an UR tool with a KR tool. On a different side, the possible world account given here to DS theory is strongly related to other possible world based accounts given to DS theory or to probability theory (e.g. Gaifman, 1986; Ruspini, 1986; Fagin and Halpern, 1989). However, our focus is the decomposition of the uncertain knowledge representation task in its KR and UR components, using intensions as a formal bridge: possible worlds are just one of the possible choices for representing intensions. Moreover, while the accounts above are normally restricted to the propositional case, we attach it to an arbitrary KR language.

[10] This suggestion is similar to some recent proposal to use an ATMS for implementing DS theory (e.g. Laskey & Lehner, 1989).

Acknowledgements. This research has benefit from worthy discussions with (and comments from) Yen-Teh Hsia, Robert Kennes, Bruno Marchal, Philippe Smets, and Nic Wilson. Fabrizio Sebastiani has an early responsibility in my all-hybridizing impetus.

References

Brachman, R.J. and Levesque, H.J. (1982) "Competence in Knowledge Representation", Proc. of AAAI-82: 189-192.
Brachman, R.J. and Levesque, H.J. (Eds.) (1985) *Readings in Knowledge Representation* (Morgan Kaufmann, CA).
Brachman, R.J., Pigman Gilbert, V., Levesque, H.J. (1985) "An Essential Hybrid Reasoning System: Knowledge and Symbol Level Accounts of Krypton", Proc. of IJCAI-85.
de Kleer, J. (1986) "An Assumption-Based Truth Maintenance System" *Artificial Intelligence* **28**: 127-162.
Fagin, R. and Halpern, J.Y. (1989) "Uncertainty, Belief and Probability", Proc. of IJCAI-89: 1161-1167.
Gaifman, H. (1986) "A Theory of Higher Order Probabilities", Proc. of the Conf. of Theoretical Aspects of Reasoning about Knowledge: 275-292.
Halpern, J.Y. and Moses, Y (1985) "A Guide to the Modal Logics of Knowledge and Belief", Proc. of IJCAI-85.
Henkind, S.J. and Harrison, M.C. (1988) "An Analysis of Four Uncertainty Calculi", *IEEE Trans. on Systems, Man and Cybernetics* **18**(5): 700-714.
Israel, D.J. and Brachman, R.J. (1981) "Distinctions and Confusions: a Catalogue Raisonné", Proc. of IJCAI-81.
Laskey, K.B. and Lehner, P.E. (1989) "Assumptions, Beliefs and Probabilities", *Artificial Intelligence* **41**(1): 65-77.
Levesque, H.J. (1984) "Foundations of a Functional Approach to Knowledge Representation", *Artif. Intell.* **23**.
Ruspini, E.H. (1986) "The Logical Foundations of Evidential Reasoning", Tech. Note 408, SRI Int. (Menlo Park, CA).
Saffiotti, A. (1987) "An AI view of the treatment of uncertainty", *The Knowledge Engineering Review* **2**(2).
Saffiotti, A. (1990) "A Hybrid Belief System for Doubtful Agents", to appear in: Procs. of the Third IPMU Conference (Paris, France, 2-6 July 1990).
Saffiotti, A. and Sebastiani, F. (1988) "Dialogue Modelling in M-Krypton, a Hybrid Language for Multiple Believers", Proc. of the 4th IEEE Conf.on AI Applications: 56-60.
Schefe, P. (1980) "On Foundations of Reasoning with Uncertain Facts and Vague Concepts", *Int. J. Man-Machine Studies* **12**: 35-62.
Shafer, G. (1976) *A Mathematical Theory of Evidence* (Princeton University Press, Princeton).
Smets, Ph. (1988) "Belief Functions", in: Smets, Mamdani, Dubois, and Prade (Eds.) *Non-Standard Logics for Automated Reasoning* (Academic Press, London).
Thompson, T.R. (1985) "Parallel Formulation of Evidential Reasoning Theories", Procs. of IJCAI-85, 321-327.
Zadeh, L.A. (1989) "Knowledge Representation in Fuzzy Logic", *IEEE Trans. on Knowledge and Data Eng.* **1**(1).