Published in K. Morik (ed.), *GWAI-87, 11th German Workshop on Artificial Intelligence*, p. 114–123,
Geseke, Germany, September 1987. Springer-Verlag

Issues of Integration and Balancing in Hybrid Knowledge Representation Systems*

Bernhard Nebel, Kai von Luck
Technische Universtät Berlin
Sekr. FR 5–8
Franklinstraße 28/29
D-1000 Berlin 10

Abstract

In the last several years the hybrid approach to Knowledge Representation has received much attention, because it was felt that one monolithic knowledge representation formalism cannot meet all representational demands. In this paper we will present one particular hybrid knowledge representation system, BACK, concentrating on matters of how to integrate different subformalisms and their interpretation. In particular, 'balancing the expressiveness' of the respective subformalisms and combining the reasoning of the subsystems in a sound way is discussed. This will lead to a new view on the realization inference, first described by Mark, as a process of constraint propagation.

1 Introduction

Aaron Sloman pointed out very clearly in [Sloman 85] the need for different knowledge representation formalisms for the adequate representation of a realistic portion of the world. This position contrasting the view that one uniform formalism is sufficient (e.g., [Kowalski 80]) is nowadays widely accepted. The most visible consequence of this is the emerging number of systems—so-called *hybrid systems*—supporting representation of knowledge by more than one formalism. The actual situation in research on hybrid systems is described in [Brachman et al 85, p. 532]:

> Many of the today's knowledge representation systems offer their users a choice of more than one language for expression of domain knowledge. While the idea has been important to the field for many years, 'multiple representations' seems to have recently become a popular catch phrase. Many of the modern expert system development environments wave the polyglot banner, and except perhaps for some stalwart first-order logicians, most everyone would probably agree that one uniform language will not serve all representational needs.
>
> It is sometimes difficult to discern the true value of multiple languages; some of the commercial development tools seem simply to appeal to the 'the more the merrier', without any clear idea of how merrier is better.

First of all, let's clarify some terms. If we talk about a *knowledge representation formalism*, we mean a formal language with given *syntax* and *semantics*[1]. This language can be analyzed for its

*This work was partially supported by the EEC and is part of the ESPRIT project 311, which involves the following participants: Nixdorf, Olivetti, Bull, Technische Universität Berlin, Universita di Bologna, Universität Hildesheim and Universita di Torino.

[1]In the sense of, e.g., [McDermott 78] or [Hayes 79].

expressivness, naturalness with respect to its intended purpose, decidability and inherent complexity of inference algorithms with respect to the formal semantics.

A *knowledge representation system* is a 'materialization' of a formalism supporting the *interpretation* of well-formed expressions of a knowledge representation language by means of inference algorithms realizing the semantics to a certain extent. These inference algorithms should be *sound*—they should not lead to wrong propositions—and in the ideal case they should be *complete*—be able to deduce any true proposition. As it turns out, however, in the real world the latter goal can only be achieved if either the formalism is very simple or if we allow for arbitrary long computations. Therefore, often the solution is to provide only the obvious, easy to compute inferences and ignore the difficult ones.

Of course, a knowledge representation system is more than just a mechanized reasoner. Another task for such a system is the maintenance of represented knowledge, i.e., it has to account for *additions* to and *updates* of the represented knowledge. In fact, this is the main difference between a knowledge representation system and a static deductive calculus (or a programming language). Furthermore, any knowledge representation system claiming to be usable has to provide a friendly interface to the human user[2]. In the sequel, however, we will ignore the latter and focus on the former two points.

In the ideal case a hybrid knowledge representation system is an implementation of a hybrid knowledge representation formalism, consisting of two or more different subformalisms. However, the mere combination of formalisms does not necessarily result in a hybrid formalism. A kind of 'glue' is needed in order to constitute a hybrid formalism consisting of

- a *representational theory* (explaining what knowledge is to be represented by what formalism) and
- a *common semantics* for the overall formalism (explaining in a semantic sound manner the relationship between expressions of different subformalisms).

The representational theory should explain why there are different subformalism, what their benefits are, and how they relate to each other. An answer should at least refer to adequacy criteria such as (cf. [McCarthy, Hayes 69]):

- *epistemological adequacy*, i.e., that the subformalisms are necessary to represent epistemological different kind of knowledge (e.g., analytic and contingent knowledge), or
- *heuristic adequacy*, i.e., that the different subformalisms permit representation of the same knowledge in different ways for reasons of efficiency.

A necessary precondition for gluing things together is that their shapes fit, a fact which might be violated in designing a hybrid formalism, at least in the case where the subformalisms are intended to represent epistemological different kinds of knowledge. For example, if one subformalism permitted definition of terms by using time relationships, but none of the other subformalisms referred to time at all, the subformalisms would be in some sense *unbalanced*. This, however, can be easily uncovered by inspecting the common semantics.

There is a more subtle sense of how two formalisms in such a hybrid system can be unbalanced, which has to do with the fact that most knowledge representation systems are necessarily *incomplete* in their reasoning in order to provide answers in reasonable time. Because of this incompleteness there could be situations where one subformalism allows to express something which obviously should have some impact on another subformalism according to the common semantics, but the system does not realize this because it's reasoning is incomplete in this aspect. This 'black hole' might be there because the incompleteness has principal reasons or because the subformalism is not heuristically adequate for this aspect. In any way, the subformalisms of the system appear to be unbalanced.

[2]And usually, a large fraction of code in a knowledge representation system is devoted to the user interface.

Although the term *balancedness* is a little bit vague, it can be captured by the following *principle of balancedness in hybrid representation systems*:

> If a representation construct in a subformalism of a hybrid formalism suggests that its usage has some impact on knowledge represented with another subformalism (according to the common semantics), then this should be realized by the system.

An example for a system with unbalanced subformalisms is KL-TWO [Vilain 85]: While it is possible to define concepts with a very rich language, only a fraction of it is used for stating contingent propositions. In particular, the number restrictions used in the NIKL subformalism has only a very limited impact on the PENNI subformalism, because the latter is not heuristically adequate to deal with cardinalities.

In the last few years, much effort has been devoted to the development of hybrid systems. Besides systems favoring multiple representations for the sake of naturalness and efficiency of the represented knowledge, systems combining formalisms for the representation of knowledge according to the distinction Frege made between meaning (*Sinn*) and reference (*Bedeutung*) were developed. In particular, the connection of KL-ONE [Brachman, Schmolze 85] derivates as formalisms for representing terminological knowledge (*TBoxes*) with formalisms for representing assertions about the actual state of the world (*ABoxes*) has been investigated[3] (e.g., KRYPTON [Brachman et al 85], KL-TWO [Vilain 85], MESON [Edelmann, Owsnicki 86], BACK [Luck et al 87] and KANDOR [Patel-Schneider 84]). The main points in this research were the design of an appropriate ABox, sometimes requiring a restriction of the TBox (e.g., in KRYPTON), and developing means for connecting the reasoning of TBox and ABox.

In the sequel we will present one particular solution to these problems pursued in BACK[4]. The following design criteria have been taken into account in developing BACK:

- The subformalism of the BACK system should be balanced.
- The BACK formalism should permit tractable inference algorithms covering almost all possible inferences[5].
- The ABox formalism of BACK should be able to represent incomplete knowledge in a limited manner (cf. [Luck et al 86]).
- The BACK system should allow for extending the knowledge base incrementally (we do not consider retractions!).
- The BACK system should reject ABox entries which are inconsistent.

The rest of the paper is divided in four parts. The next section gives a very brief introduction to hybrid KL-ONE systems for those readers unfamiliar with this topic. Then the BACK ABox is discussed, and it is shown how we achieved the goals concerning the formalism stated above. Finally, we investigate what kind of inferences the combination of the ABox and the TBox permit, and how these inferences can be realized. This leads to a new view on the *realization* inference as discussed in [Mark 82], which can be characterized as a 'constraint propagation' process.

---

[3] A rough similarity of this distinction is known in the database area with the distinction of database schemata and database contents, but these approaches take the schema definition as a source of integrity constraints and not as a formalism for the intensional definitions of terms.

[4] The **B**erlin **A**dvanced **C**omputational **K**nowledge Representation System.

[5] Unfortunately, complete and tractable inference algorithms are possible only for very simple TBoxes (cf. [Nebel 87] and [Brachman, Levesque 84]).

2 A Short Characterization of Hybrid KL-ONE Systems

KL-ONE as described in [Brachman, Schmolze 85] is perfectly well suited for the introduction of a terminology. It allows to specify *concepts* by stating superconcepts and restrictions on relationships to other concepts, which are called *roles*. Concepts and roles correspond roughly to *generic frames* and *slots* in the frame terminology. Unlike frames, however, concepts are understood as purely *intensional*. Furthermore, concepts may be introduced as *primitive* or *defined*. In the former case, the concept description specifies only the *necessary* conditions, while in the latter case the concept description is *necessary* and *sufficient*.

The following is an (informal) example for the definition of the (slightly artificial) concept modern-small-team. Assuming that human, woman, and team are already introduced as concepts and that member is a role with leader as a *subrole*, the definition might look as follows:

A modern-small-team is (defined as)
a team and
having at most 4 members and
all members are humans and
having exactly 1 leader and
all leaders are woman

This example shows some of the important *concept-forming operators* which may be used in order to create new concepts, namely *specialization*, *number restriction* and *value restriction* of a role. In KL-ONE, concepts and roles are usually presented by graphical means (cf. Fig. 1) which may be sometimes easier to comprehend. However, it does not make any difference whether a linear notation or a graphical network is used provided that a formal semantics can be given (cf. [Luck et al 87], [Luck, Owsnicki 87]).

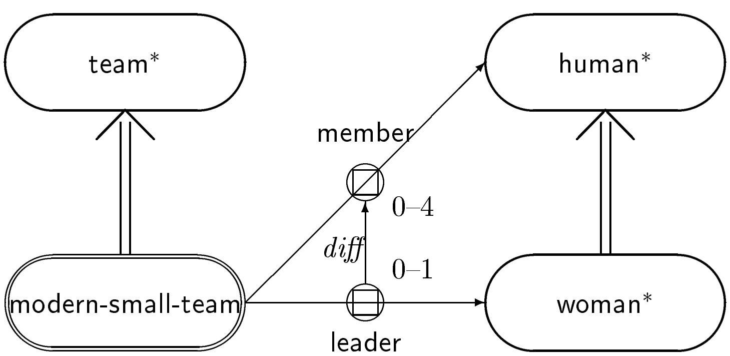

Figure 1: An example of a KL-ONE concept

While KL-ONE does very well in the area of defining concepts, it is weak in stating any contingent truth about the world, for example that "KIM is a woman, and that she is a member of some modern-small-team, but not the leader of that team". Another example would be that "there is at least one modern-small-team". In order to compensate for this, an ABox has to be employed, a topic we will discuss below.

3 The Design of the BACK ABox

The formalisms for representing terminological knowledge in the hybrid systems mentioned above differ only in what concept- and role-forming operators they provide—sometimes motivated by arguments concerning the computational complexity. The situation with ABoxes, however, is quite different. Here we meet a variety of approaches, e.g.:

KRYPTON Full first-order predicate logic.

KL-TWO Variable-free predicate logic with equality.

KANDOR Object-centered, frame-like schema.

MESON Object-centered, frame-like schema.

Using predicate logic (first-order or restricted) as the ABox has several advantages. One is that it gives a clear account to a *common semantics* for the entire formalism[6]. Another benefit of predicate logic is its *plasticity* as Hayes called it [Hayes 74]—the possibility to give *partial descriptions* and to extend a knowledge base *incrementally*. On the other hand, there are a lot of disadvantages. One obvious problem with unrestricted first-order predicate logic are the computional costs, a fact which exclude such systems from being used as a knowledge representation system in practical AI systems. Besides that, there are also problems with the expressiveness of predicate logic. For instance, there is no easy way to state that a given description is *exhaustive* or that different constants denote different objects. The latter two properties are responsible for the fact that *number restrictions* on roles are omitted from the KRYPTON TBox and that they are almost mere comments in the KL-TWO TBox.

For frame-like ABoxes as employed in KANDOR or MESON the situation is opposite. The *unique name hypothesis* and *closed world assumption* are taken for granted, very similar to conventional data bases. This forces the user of such a system to give a *complete* description of the state of affairs and prohibits incremental 'monotonic' additions. However, this kind of representation—called 'vivid' by Levesque in [Levesque 86]—allows for elegant and simple inference algorithms. In the case of hybrid KL-ONE systems it implies that number restrictions and value restrictions of the TBox can be fully utilized in the ABox. A short note about formal semantics might be in order here. While frame-like schemata are often viewed as some ad hoc data structure, it is nevertheless possible to specify a precise formal semantics for at least a subset (cf. [Hayes 79]).

In BACK we tried to combine the benefits of both approaches. An object-centered language was chosen to describe objects of the domain. However, instead of insisting that all role fillers for a given object have to be specified, the information can be *incomplete* and may be refined later. For this purpose the follwing operators are provided:

- stating the cardinality of a role-filler set, e.g., "MARY has at least 2 children";
- disjunctive information on role fillers, as e.g., "MARY is married to JOHN or TOM";
- stating the exhaustiveness of the provided information selectively (cf. [Hayes 74]), e.g., "MARY has as a friend TOM, and these are all friends MARY has".

In [Luck 86][7] it was shown that these operators are sufficient to achieve the goal of 'balanced expressiveness' at least with respect to number restrictions. This also applies to the other concept-forming operators introduced in section 2 as we will see below. Furthermore, the formalism permits the representation of *incomplete knowledge*, *incremental refinement* as well as the treatment of *complete information*. Finally, because of the careful selection of ABox operators, reasoning in the ABox appears to be *tractable* at least in a weak sense.

4 Combining the Reasoning of ABox and TBox

While it appears to be very obvious what kind of inferences are granted by the separate formalisms, the interesting problems stand up if we combine the reasoning of the TBox and the ABox.

[6] In these approaches concepts are viewed as 1-place predicates and roles as 2-place predicates.

[7] This article also gives a more formal account of the ABox semantics we ommited here because of lack of space.

The TBox alone is only good for *terminological inferences*, including detecting implicit specalization relationships between concepts, recognizing contradictory concepts and computing the properties inherited from superconcepts. The ABox alone, depending on its flavor, can draw the inferences sanctioned by first-order logic (in the KRYPTON case), by variable-free predicate logic with equality (in KL-TWO), or it can infer whether a domain object is related to other objects, and how many there are (in the case of object-centered ABoxes).

In order to demonstrate the interaction between ABox and TBox in BACK let us investigate a small (informal) example. Let us assume the following concept definitions:

A man is a
    human.
A woman is a
    human.
man and woman are disjoint concepts.
A team is (defined as) a
    set with members which are all humans, and
    any leaders are members as well.
A male-team is (defined as) a
    team with all members are men.
A small-team is (defined as) a
    team with at most 4 members.
A modern-small-team is (defined as) a
    small-team with at least 1 leader and
    all leaders are women.

Now let us assume that the following object decriptions are given:

TOM, DICK and HARRY are instances of
    the concept man.
JUNK is an instance of
    modern-small-team with
        TOM, DICK, HARRY and KIM
        as members of the team.
CHAUVIS is an instance of
    male-team with
        KIM
        as one of its members.

After this sequence, nothing is wrong with the contents of the knowledge base at first sight. If, however, the situation is more thorougly analyzed, it becomes obvious that

- after the JUNK team is introduced, KIM is known to be human, because of the value restriction on member;
- after the CHAUVIS team is introduced, it becomes clear that KIM is a man with the same argument;
- and because of that, nobody can become the leader of the JUNK team, because the leader subrole can only be filled with a woman.

Thus, the ABox contents is *contradictory*. The JUNK team cannot exist in this configuration and simultaneously be called a modern-small-team, or otherwise, KIM cannot a member of the male-team.

While the arguments above are only informal, it is, of course, possible to formalize the problem using the common semantics of the formalism [Luck et al 87] and deduce a contradiction as follows:

$\forall x\colon \mathsf{man}(x) \quad \Rightarrow \quad \mathsf{human}(x)$ (1)

$\forall x\colon \mathsf{woman}(x) \quad \Rightarrow \quad \mathsf{human}(x)$ (2)

$\forall x\colon \neg\mathsf{man}(x) \vee \neg\mathsf{woman}(x)$ (3)

$\forall x\colon \mathsf{team}(x) \quad \Leftrightarrow \quad \mathsf{set}(x) \wedge (\forall y\colon \mathsf{member}(x,y) \Rightarrow \mathsf{human}(y)) \wedge$
$(\forall z\colon \mathsf{leader}(x,z) \Rightarrow \mathsf{member}(x,z))$ (4)

$\forall x\colon \mathsf{male\text{-}team}(x) \quad \Leftrightarrow \quad \mathsf{team}(x) \wedge (\forall y\colon \mathsf{member}(x,y) \Rightarrow \mathsf{man}(y))$ (5)

$\forall x\colon \mathsf{small\text{-}team}(x) \quad \Leftrightarrow \quad \mathsf{team}(x) \wedge$
$(\exists\, y_1, y_2, y_3, y_4\colon \mathsf{member}(x,y_1) \wedge \mathsf{member}(x,y_2) \wedge$
$\mathsf{member}(x,y_3) \wedge \mathsf{member}(x,y_4) \wedge$
$(\forall z\colon \mathsf{member}(x,z) \Rightarrow z = y_1 \vee z = y_2 \vee z = y_3 \vee z = y4))$ (6)

$\forall x\colon \mathsf{modern\text{-}small\text{-}team}(x) \Leftrightarrow \quad \mathsf{small\text{-}team}(x) \wedge (\exists z : \mathsf{leader}(x,z)) \wedge$
$(\forall y\colon \mathsf{leader}(x,y) \Rightarrow \mathsf{woman}(y))$ (7)

$\mathsf{TOM}{\neq}\mathsf{DICK} \wedge \mathsf{TOM}{\neq}\mathsf{HARRY} \wedge \mathsf{TOM}{\neq}\mathsf{KIM} \wedge \mathsf{DICK}{\neq}\mathsf{HARRY} \wedge \mathsf{DICK}{\neq}\mathsf{KIM} \wedge \mathsf{HARRY}{\neq}\mathsf{KIM}$[8] (8)

$\mathsf{man}(\mathsf{TOM}) \wedge \mathsf{man}(\mathsf{DICK}) \wedge \mathsf{man}(\mathsf{HARRY})$ (9)

$\mathsf{modern\text{-}small\text{-}team}(\mathsf{JUNK}) \wedge \mathsf{member}(\mathsf{JUNK},\mathsf{TOM}) \wedge \mathsf{member}(\mathsf{JUNK},\mathsf{DICK}) \wedge$
$\mathsf{member}(\mathsf{JUNK},\mathsf{HARRY}) \wedge \mathsf{member}(\mathsf{JUNK},\mathsf{KIM})$ (10)

$\mathsf{male\text{-}team}(\mathsf{CHAUVIS}) \wedge \mathsf{member}(\mathsf{CHAUVIS},\mathsf{KIM})$ (11)

by (5,11): $\mathsf{man}(\mathsf{KIM})$ (12)

by (3,9,12): $\neg\mathsf{woman}(\mathsf{KIM}) \wedge \neg\mathsf{woman}(\mathsf{TOM}) \wedge \neg\mathsf{woman}(\mathsf{DICK}) \wedge \neg\mathsf{woman}(\mathsf{HARRY})$ (13)

by (6,7,8): $\forall z : \mathsf{member}(\mathsf{JUNK},z) \Rightarrow \mathsf{z{=}TOM} \vee \mathsf{z{=}DICK} \vee \mathsf{z{=}HARRY} \vee \mathsf{z{=}KIM}$ (14)

by (4,7): $\exists z : \mathsf{member}(\mathsf{JUNK},z) \wedge \mathsf{woman}(z)$ (15)

by (14,15): $\mathsf{woman}(\mathsf{KIM}) \vee \mathsf{woman}(\mathsf{TOM}) \vee \mathsf{woman}(\mathsf{DICK}) \vee \mathsf{woman}(\mathsf{HARRY})$ (16)

by (13,16): *contradiction*

5 Realization as Constraint Propagation

In order to find such contradictions and, more generally, to determine the concept which most accurately describes a given object a forward inference technique called *realization* [Mark 82] is usually employed. Realization is very similar to *classification*, an inference technique used to maintain the taxonomy of concepts in the TBox [Lipkis 82]. In fact, realization can be viewed as *abstraction*—generating a description in terms of the TBox—followed by *classification* of this description (cf. [Vilain 85]).

A first approximation to the implementation of this inference could be realized as follows. After a new assertion about an indivdual enters the ABox (either a new concept or a new role filler) the following has to be done:

- Propagate all role fillers of subroles to the corresponding superrole.
- Determine the cardinality of the role filler sets for each role (this can be a range in the case of incomplete information). This cardinality information is used as the actual number restriction for the role in the abstraction process.
- If a role-filler set is closed, i.e., all potential candidates are known, then the *generalization*[9] of the descriptions of all potential role fillers can serve as the actual value restriction in the abstraction process.
- Now, the generated number and value restrictions, the old description of the individual as well as the new one can be used to construct a concept definition which can be classified and after that serve as the new *most general specialization* (MSG) of the individual under investigation.

[8] This statement expresses the *unique name hypothesis*.

[9] Generalization is not a concept-forming operator and there seems to be no easy way to assign a compositional semantics to it. However, it is easy to build such a concept structurally, which is the most specialized one subsuming a given set of concepts.

This algorithm does take into account all information which is supplied *locally* to an individual, but ignores any non-local consequences. In the example of the last section, we would at least require that after the abstraction process the role fillers are to be specialized according to the value restriction of the male-team. Additionally, we note that in order to detect the contradiction in the example we also have to account for the case that the specialization of an individual (KIM from human to man) can lead to the specialization of another individual (JUNK from modern-small-team to the empty concept) the first one is a role filler of.

If we analyze this algorithm more thoroughly, we may note that this process can trigger other specializations which in turn may propagate restrictions. Because this sounds very expensive from a computational point of view, one could argue that because the system reasoning process is incomplete anyway, it is legitimate to restrict the resources allocated for the realization process and leave such situations alone. And this was indeed the first approximation to a solution we chose. However, this is not a general incompleteness of the inference algorithm, but it depends on the order of input! If, in our example, the order of input between the JUNK and the CHAUVI would have been reversed the contradiction would be easily detected.

That is certainly not the kind of behaviour we expect from a knowledge representation system. Although it is clear that we have to live with incomplete reasoners, this incompleteness should be systematic, perhaps even describable by a systematic, model-theoretic approach (cf. [Patel-Schneider 86]). In conclusion, if we claim to *integrate* the reasoning of two formalisms, a minimal requirement is that the inferences are independent of the order of input.

The only solution to this problem is to employ some sort of 'constraint propagation'[10]. After a new MSG is determined for some individual I,

- all role fillers at I have to be specialized to the corresponding value restriction (forward propagation);
- additionally, all individuals, which mention I as a role filler have to be checked whether it is possible to specalize them (backward propagation); this can only happen if the corresponding role-filler set is closed;
- finally, these steps have to be performed for all individuals which get new MSGs during this process.

For the example above, this suffices to detect the contradiction. The JUNK team will be specialized to the following concept:

> a modern-small-team and
> all of its members are men

In classifying this concept, the value restriction of the member role is propagated to the subrole leader (a subrole has to adhere to the value restriction of its superrole) resulting in a value restriction of (AND man woman) which apparently is contradictory.

Altogether, this process assures that incompleteness depending on the order of input cannot appear. However, it also sounds very costly from a computational point of view. It is even not very clear whether this process always terminates.

Fortunately, the process of propagating MSG is not a general constraint propagation process. A first fact about this process, we may note, is that backtracking cannot occur. That means if we encounter a contradiction there is no way to resolve it, but to reject the input which led to the contradiction. This tells us that we will not get a combinatorial explosion because of reasoning by case. There are more facts which constrain the propagation space:

[10]While we first thought that BACK is the only system using this technique, because the literature does not give any hints in this direction, latter on we learned from Marc Vilain that KL-TWO works in a similar way.

- An individual which gets a new MSG by backward propagation cannot trigger new forward propagations. This is because the new MSG does reflect the current role fillers as necessary and sufficient conditions, i.e., there cannot be any new restrictions on role fillers.

- Backward propagation leads to a new MSG, if the role-filler set is closed and all other role fillers are already more specialized than the value restriction, a situation which does not occur very often.

- Forward propagation has a significant non-local effect only if the chain of individuals the value restrictions are propagated along corresponds to a line of concepts which are specialized in parallel, which is not very likely.

- In the worst case, the number of recomputations of an MSG during forward and backward propagation is bounded by the product of the number of individuals and the number of concepts in the TBox (before the entire process starts). This, however, would result in MSGs which cannot be specialized further, i.e., the upper bound can be divided by the number of input operations.

Finally, we should discuss where incompleteness might arise in in the integrated reasoning process. First, realization cannot be more complete than classification. This means that some very weird cases of contradictions cannot be recognized (cf. [Nebel 87]). Furthermore, the abstraction process is a source of incompleteness as well. In our example, we could prove that KIM is necessarily a woman after the first two inputs; however, the abstraction failed to recognize this. The reason is, that this would require some kind of reasoning by case we strictly avoided because of the computional costs.

6 Conclusion

We have argued that a hybrid knowledge representation system should be balanced in its expressiveness and integrate the reasoning of its subformalisms in a sound manner based on a common semantics. Furthermore, we made plausible that the BACK system does very well in these respects. By giving a thorough description of an inference technique known as ‘realization’, we demonstrated that this inference can only be realized as a constraint propagation process, a very special and limited one, though.

Acknowledgement

Many of the ideas presented here are the result of the work in the BACK project. The other members of this project, Christof Peltason and Albrecht Schmiedel, deserve much of the credit for these ideas. Additionally, the discussions with Marc Vilain and Bob MacGregor were very helpful in shedding some light on the realization process.

References

[Brachman, Levesque 84] Ronald J. Brachman, Hector J. Levesque, The tractability of subsumption in frame-based description languages, in: *Proc. AAAI-84*, Austin (Tex.), August 1984, 34–37.

[Brachman, Schmolze 85] Ronald J. Brachman, James G. Schmolze, An overview of the KL-ONE knowledge representation system, in: *Cognitive Science*, **9**(2), April–June 1985, 171–216.

[Brachman et al 85] Ronald J. Brachman, Victoria Pigman Gilbert, Hector J. Levesque, An Essential Hybrid Reasoning System: Knowledge and Symbol Level Accounts in KRYPTON, in: *Proc. 9th IJCAI*, Los Angeles (Cal.), 1985, 532–539.

[Edelmann, Owsnicki 86] Jürgen Edelmann, Bernd Owsnicki, Data Models in Knowledge Representation Systems: A Case Study, in: C.-R. Rollinger, W. Horn (eds.), *GWAI-86 und 2. Österreichische Artificial-Intelligence-Tagung*, Springer, Berlin (Germany), 1986, 69–74.

[Hayes 74] Patrick J. Hayes, Some Problems and Non-Problems in Representation Theory, *Proc. AISB Summer Conference*, University of Sussex, 1974, 63–79.

[Hayes 79] Patrick J. Hayes, The Logic of Frames, in: D. Metzing (ed.), *Frame Conceptions and Text Understanding*, deGruyter, Berlin (Germany), 1979, 46–61.

[Kowalski 80] Robert A. Kowalski, Contribution to the SIGART Special Issue on Knowledge Representation, in: *SIGART Newsletter*, No. 70, February 1980.

[Levesque 86] Hector J. Levesque, Making Believers out of Computers, Computers and Thougth Lecture at IJCAI-85, Los Angeles (Cal.), published in: *Artificial Intelligence*, **30**(1), 1986, 81–108.

[Lipkis 82] Tom Lipkis, A KL-ONE Classifier, in: J. G. Schmolze, R. J. Brachman (eds.), *Proc. 1981 KL-ONE Workshop*, BBN Technical Report 4842, Bolt, Berank and Newman, Inc., Cambridge (Mass.), 1982, 128–145.

[Luck 86] Kai von Luck, Semantic Networks with Number Restricted Roles or Another Story about Clyde, in: C.-R. Rollinger, W. Horn (eds.), *GWAI-86 und 2. Österreichische Artificial-Intelligence-Tagung*, Springer, Berlin (Germany), 1986, 58–68.

[Luck, Owsnicki-Klewe 87] Kai von Luck, Bernd Owsnicki-Klewe, Neuere KI-Formalismen zur Repräsentation von Wissen, to appear in: T. Christaller (ed.), *Künstliche Intelligenz*, Springer, Berlin (Germany), 1987.

[Luck et al 86] Kai von Luck, Bernhard Nebel, Christof Peltason, Albrecht Schmiedel, BACK to Consistency and Incompleteness, in: H. Stoyan (ed.), *GWAI-85*, Springer, Berlin (Germany), 1986.

[Luck et al 87] Kai von Luck, Bernhard Nebel, Christof Peltason, Albrecht Schmiedel, *The Anatomy of the BACK System*, KIT Report 41, Technische Universität Berlin, Berlin (Germany), January 1987.

[Mark 82] William Mark, Realization, in:J. G. Schmolze, R. J. Brachman (eds.), *Proc. 1981 KL-ONE Workshop*, BBN Technical Report 4842, Bolt, Berank and Newman, Inc., Cambridge (Mass.), 1982, 78–89.

[McCarthy, Hayes 69] John McCarthy, Pat Hayes, Some Philosophical Problems from the Standpoint of Artificial Intelligence, in: B. Meltzer, D. Michie (eds.), *Machine Intelligence* **4**, Edinburgh University Press, Edinburgh (Great Britain), 1969, 463–502.

[McDermott 78] Drew McDermott, Tarskian semantics, or no notation without denotation!, in: *Cognitive Science*, **2**(3), July–September 1978, 277–282.

[Nebel 87] Bernhard Nebel, *Computational Complexity of Terminological Reasoning in BACK*, to appear as a KIT Report, Technische Universität Berlin, Berlin (Germany), 1987.

[Patel-Schneider 84] Peter F. Patel-Schneider, Small can be Beautiful in Knowledge Representation, in: *Proc. IEEE Workshop on Principles of Knowledge-Based Systems*, Denver (Colo.), 1984, 11–16.

[Patel-Schneider 86] Peter F. Patel-Schneider, A Four-Valued Semantics for Frame-Based Description Languages, in: *Proc. AAAI-86*, Philadelphia (Pa.), August 1986, 344–348.

[Sloman 85] Aaron Sloman, Why We Need Many Knowledge Representation Formalisms, in: M. Bramer (ed.), *Research and Development in Expert Systems*, Cambridge University Press, Cambridge (Great Britain), 1985.

[Vilain 85] Marc B. Vilain, The Restricted Language Architecture of a Hybrid Representation System, in: *Proc. 9th IJCAI*, Los Angeles (Cal.), 1985, 547–551.