# THE INTERACTION WITH INCOMPLETE KNOMEDGE BASES: A FORMAL TREATMENT* 

Hector J. Levesque<br>Department of Computer Science<br>University of Toronto<br>Toronto, Ontario


#### Abstract

Some formal representation Issues underlying the Interaction between an expert system and Its knowledge base are discussed. It Is argued that a language that can refer both to the application domain and to the state of the knowledge base Is required to specify and to question an incomplete knowledge base. A formal logical language with this ability is presented and its semantics and proof theory are defined. It is then shown how this language must be used to interact with the knowledge base.


## 1. INTRODUCTION

An Important characteristic of any knowledgebased system Is its interaction with a knowledge base (or KB for short) that provides and maintains information about the application domain. In general, there are two distinct modes of interaction between an expert system and its KB: the system will want to ask and to tell the KB about the application area. In this paper, I will examine what special expressive requirements are placed on the language(s) of interaction when a system must deal with a KB that is Incomplete.

Loosely speaking* a KB is Incomplete when it does not have all the Information necessary to answer some question of Interest to the system. When a knowledge-based system is forced to depend on an incomplete KB, its ability to make decisions or solve problems is seriously compromised. In some cases, the lack of knowledge can be circumvented using general defaults [1], while in other situations, special heuristics are required [2]. However, no matter how a system plans to deal with incompleteness, it must first be able to determine where this incompleteness lies. In other words, a system has to find out exactly where knowledge is lacking before it can decide what to do about it. This suggests that a $K B$ must be capable of providing information not only about the application area, but about Itself as well. Thus, the language used to interact with a KB must allow a user to define and inquire about what the KB does and does not know. The major issue of this paper, then, is what impact this capability will have on the language of interaction.

[^0]The approach taken in my work [3] focuses primarily on the formal aspects of this issue. As described in section two, l will be dealing with KB's consisting of formulas of the first order predicate calculus. In section three, I examine what questions one would like to be able to ask such a KB and conclude that these questions are best phrased in a superset of the predicate calculus I call KL. In section four, I examine the language KL in detail and provide a semantics and proof theory. Given this analysis of KL, I then consider its use as a language to interact with a KB: in section five, the formulas of KL are used to query a KB and I show that this involves a non-monotonic version of the language; in section six, the formulas of KL are used to define a $K B$ and I indicate why this requires converting the formulas into the language of the KB.

Three areas of related research should be noted. First of all, my formalization of KL owes much to the work on the logic of knowledge and belief pioneered in [4] and continued in (5] and [6] (though my application is somewhat different). Secondly, there has been research on incompleteness from the database area ([7],[8]) concentrating primarily on efficient query evaluation. Finally, although the research here does not deal with defaults or assumptions, the work on non-monotonic reasoning ([9],(10],[11], [12]) has provided technical inspiration.

## 2. FIRST ORDER KNOMFDGE BASES

As described in [13], a first order $K B$ is a finite set of closed formulas of first order logic (FOL). I will restrict my attention to consistent sets of formulas that do not contain function symbols other than a (possibly infinite) set of O-ary symbols called constants. Moreover, a bijection is assumed to exist between the set $C$ of constants and the set $D$ of entitles in the domain of discourse.

A query in this framework is any formula of FOL and is answered by consulting the provability relation $H$ to determine what does and does not follow from the KB. For each closed query there are three possible replies: yes, no or unknown, depending on whether the query, its negation or neither follow from what is available in the KB. A KB is incomplete when some query has an unknown answer and complete otherwise.

Consider a typical KB containing facts such as

Teachea(John, bill), Student(bill),...
To avoid having to include within this $K B$ all the "negative" facts about the application such as
-Student (John), nTeaches(bill,john),..•
an assumption is usually made (once and for all) that the negation of any atomic formula can be inferred from the inability to infer the atomic formula. This is the closed world assumption (CWA) [14] which, along with the assumption that there is a bljection between constants and entitles in the domain, guarantees that any KB is complete.

If, on the other hand, we have a $K B$ that contains
[Student(mary) v Student(susan)]
but does not contain either
Student(mary) or Student(susan),
the OWA cannot be used since it would allow us to conclude that neither Mary nor Susan are students, contradicting what is in the KB. This KB only partially describes a world, in that it specifies that one of Mary or Susan is a student but not which. Thus, there is a query that is true in the intended application but whose truth cannot be determined on the basis of what is available in the $K B$. The KB is, consequently, incomplete.

For the rest of this paper, I will say that a KB knows a formula if it answers yes to the query and, for any predicate $p$ and constant $c$, that $c$ is a known p when the formula $\mathrm{p}(\mathrm{c})$ is known.

## 3. THE RNOMFDGE BASE QUERY LANGUAGE

Even though a KB is incomplete, there may be completeness in certain areas. For example, suppose that

$$
\text { KB } \quad(x)[\operatorname{Teacher}(x)=(x-J o h n ~ v x-J i m)]
$$

In this case, not only are John and Jim teachers, but the KB knows that they are the only teachers. It therefore has a complete picture of the teachers. On the other hand, it may be the case that

## KB h (Ex)Teacher(x)

without there being any known teachers. In this situation, the KB knows that there is a teacher but does not know who and so does not have a complete list of teachers. There is, moreover, a third possibility which is that it cannot be determined from what is available in the KB whether or not it has complete knowledge of the teachers. So just as the question

## Is John a teacher?

may be answered yes, no or unknown, the question
Are all the teachers known?
may also be yes, no or unknown depending only on what is in the KB.

If the purpose of a query language is to provide an accurate picture of what is and is not available in the KB, we should be able to formulate queries that ask the KB about its incompleteness. In FOL, a query that asks if there is a teacher apart from the known teachers might be expressed by

$$
(E x)[\text { Teacher }(x) \quad-»(x « c l \vee \ldots \vee x « c k)]
$$

where the $c^{f} s$ are constants ranging over all the teachers known to the KB.

There are a couple of problems with this method of asking if all the teachers are known. First of all, there could be a very large number of known teachers. In some applications and for some predicates, there may even be an Infinite number. Secondly, we have to know what these constants are in the first place to be able to formulate the query. For a large and complex KB, it could happen that only the KB has this information. One can also Imagine situations where the KB will not divulge this information for security reasons while still being willing and able to answer questions about its incompleteness.

This suggests that the KB Itself should keep track of Its Incompleteness in the same way it maintains knowledge of the application area. One possibility, for example, Is to have a predicate "Known-teacher" and to allow a KB containing the following:

Teacher(john), Known-teacher(John), [Teacher (Jim) v Teacher (dan)],
-Known-teacher (jim), -Known-teacher (dan).
The problem with this approach (and any other that Involves a direct encoding into FOL) is the management of this extended language. There is a very definite relationship between "Teacher" and "Known-teacher" that must be captured somehow. Among other things, we want that whenever

Teacher(c)
is in the $K B$, then
Known-teacher(c)
is as well. The closest we can come to expressing this is by an axiom stating that

$$
(x)[\operatorname{Teacher}(x)=\text { Known-teacher }(x)] .
$$

However, this does not work since this formula is inconsistent with the above KB. The correspondence between "Teacher" and "Known-teacher" is much more complex. In a nutshell, "Known-teacher" should not be a predicate for the simple reason that Its truth or falsity does not depend on the domain being modelled but on the model (i.e. the KB) Itself.

The solution, then, is to leave the $K B$ as is
but to extend the query language to allow questions that refer to the current state of knowledge of the KB . The query language I propose, called KL, contains all of FOL and, in addition, has formulas of the form

Ka
read as
The KB knows that a.
This leaves us with a first order KB while still allowing us to query the KB regarding its incompleteness. In particular, we have a new form of query evaluation

## ks if a

where $\mid \boldsymbol{F}$ is some (as yet to be specified) provability relation and a is any formula of KL. For example, to find out if the KB has an Incomplete list of teachers, we ask if

Similarly, while the query
(Ex)Teacher(x)
can be used to find out if there are any teachers, the query

## ( $2 x$ ) [Teacher ( $x$ ) ]

asks if the KB knows who any of them are. To be able to define the $\mid \boldsymbol{F}$ relation, we must first look at the semantics of the language KL Itself.

## 4. THE LANGUAGE KL

The query language KL has the same formation rules as $\operatorname{FOL}$ but also includes the rule

## If $a \in \mathbf{K L}$ then $K a \in \mathbf{L I}$.

There are, consequently, two kinds of formulas in KL: the first, like

## $p(c) A(B=q(x)$

will be true or false depending only on how the world is, that is, on the interpretation of the constant and predicate symbols; the second, like

## $K[p(c) \wedge(\mathbf{L}) \mathbf{q}(x)\}$

will be true or false depending only on how the KB is, that is, on what is and is not known. I will call the latter formulas pure. There are also formulas in KL that are mixtures of the two types and whose truth value depends both on the world and the KB. KL also allows for matementedte in formulas such as

$$
\mathrm{K}[(E x) \mathrm{Kp}(x)]
$$

which talk about the KB's knowledge of its own knowledge.

The semantic Interpretation of a closed formula of KL will depend on both a world description (or interpretation, in the Tarsklan sense) and a description of a KB. In general, a KB can be viewed as a partial description of $s$ world and can thus be characterised by a set of world descriptions. To make all of this more precise, first note that because of the assumed bljection between constants snd entitles in the domain, a world description need only assign a truth value to the atomic sentences. Thus, we csn define the set of world descriptions as

## $W D=[A T O N S ~+\{T, F\}]$.

KB descriptions are, then, Just non-empty sets of world descriptions:

```
MD - {(a mD | E not mepty).
```

The interpretation of any closed formula a is provided by the function $V$.

```
    V & [KL x WD * TD }->{T,F}] is defined by
V(a,w,a)=w(a) when a is atoaic,
V(-\alpha,v,m)=T Iff V (\alpha,w,a)=F,
V(avp,w,a) = T iff V(\alpha,w,m) = T or V(p,w,m) = T,
V((EX)a,v,E) = T iff for acme c, V(\alpha<x/c>,Y,a) = T,
V(ra,w,@) = T Lff for evary w'\epsilonn, V (a,w',m) - T.
```

A formula is valid when it is true with respect to every world and KB description. In the case where a is pure, I will use $V(a, m)$ to refer to the truth value of a with respect to the KB described by $m$.

Turning now to the proof theory for KL, since the language includes FOL as a subset, we will need the axioms of FOL and its two inference rules: modus ponens snd universal generalization. To account for the K operator, we have to realize that it behaves like s provability relation in that something is known when it "follows" from what is available in the KB. We will, therefore, insist that the axioms of FOL are known and that the KB Is able to perform modus ponens and universal generalization based on what is known. This might be called the assumption of competence. As for meta-knowledge, it is convenient to assume that the KB knows the correct truth value of any pure formula. In other words, there is never any reason to tell the KB about itself nor is there any reason to doubt what the KB knows about Itself. No matter how incomplete or inaccurate a KB can be about the world, it is assumed to be the final authority on itself. This might be called the assumption of closure. Note that this assumption applies only to pure sentences. Fortunately, these can be given a syntactic characterization: a formula is pure iff every occurrence of a predicate symbol appears within the scope of a K operator. We therefore have the following axlomatization of the language KL:

```
Axtion Sche:ats
1. The axcoms of YOL
2. Ka, wherea fo an axion of FOL
3. K(\alpha>\rho) > (Xavk\rho)
4. (x)Ma}=K(x)
5. a \equivka where a lo pure
```

Modus Ponens and Universal Generalization
If we let $\boldsymbol{F}$ denote the provability relation for this axiomatization, then the key result here is that the proof theory is both sound and complete with respect to the semantics given above:

Theorem: $\boldsymbol{F} \boldsymbol{f}$ iffads valid.
One thing to notice is that by the closure property, we have that for any pure a

## 1 K Kana.

This is, however, not the case for arbitrary a in that, for example, the set

## $\left\{\rightarrow p(c), X_{p}(c)\right\}$

is satisflable and, hence, logically consistent. This is a situation where the $K B$ is behaving properly but just happens to be mistaken about the world. So, in some sense, the $K$ operator should be read as "believe" rather than "know". On the other hand, the kind of belief Involved here is very special in that a KB will always think it is dealing with knowledge since

## $1 \mathbf{K}(\mathrm{~K} \alpha=\alpha)$

for any formula a. So there is an aspect of commitment to what is believed in that a KB will never believe it has mistaken beliefs.

## 5. QUERIES REVISITED

Having examined the semantics and proof theory of KL , we are still faced with the problem of specifying what is meant by

K 8 IF $\alpha$ where K $\leq$ FOL and $\alpha \in K L$.
The idea here is that this should mean
"If 1 know only whet is in KB, then $I$ know $0 . "$
Semantlcally, this is saying that $K a$ is true with respect to the knowledge base that has the least amount of world knowledge consistent with knowing everything in KB. Thinking of a knowledge base as described by a set of world descriptions, we want the least amount of world knowledge and, thus, the largest possible set.

## Let $M(K B)=\{v \mid v a a t i a f i e s ~ t h e ~ K B\}$.

Any knowledge base that knows more than what is in the $K B$ is described by a subset of $M(J C B)$ and, so, has a more refined view of the world. This suggests how query evaluation can be defined semantlcally:

## 

So, the answer to a question is yes exactly when it is known to be true in $\mathrm{M}(\mathrm{KB})$. For example, if

KB1 |f Tancher (john),
KB1 |F-X\{Student(b11I)],
KE1 |F (Ex)K[Tescher ( $x$ )],
KB1 |FK-K[Studant(b111)].
KB1 | $\mid$ - $K[(E x)[$ Teacher ( $x$ ) A KTeacher ( $x$ )]].
The last statement above confirms that KB1 knows that it does not know if it has a complete list of teachers. Note that 11- is a non-monotonlc operator in that it is not the case that

## $K B 1+S t u d e n t(b 111) \mid f-\operatorname{KlSt}$ Stent(bili)].

However, unlike a fully general non-monotonlc logic as presented in [9] or [10], the $\mid \vdash$ operator here has been given a simple and natural semantic characterization. Below l will present a proof theoretic analogue of this operator and claim a soundness and completeness type result, but before doing so it is necessary to examine how KL can be used to specify a KB.

## 6. THE KNOMFDGE BASE DEFINITION LANGUAGE

Since, for incomplete KB's, the CWA is not used uniformly, It would be extremely convenient to be able to tell the KB when (if ever) the assumption could be used for special cases. Conversely, we should also be able to tell the $K B$ when the CNA cannot be used. If we let

## $\lambda=(\mathrm{EX})$ [Teacher $(x)$ A AKTencher (x)]

then $\boldsymbol{- \lambda}$ states that if someone is not currently known to be a teacher then he is not a teacher. So $\boldsymbol{\rightarrow} \boldsymbol{\lambda}$ is the CWA relativized to teachers while $\lambda$ itself is a statement that this assumption cannot be used (because there are teachers other than the currently known ones). The question immediately arises as to whether or not we can add formulas such as $\boldsymbol{\lambda}$ or $\boldsymbol{\square} \boldsymbol{\lambda}$ to a KB , thus generalizing a $K B$ to be any consistent set of formulas from KL instead of FOL.

There are a number of problems with this generalization but I will address only one of these relating to the formula $\lambda$. The idea here is that we would start off with a KB such as KB1 defined earlier. Since KB1 does not know whether or not John is the only teacher, it is not the case that either

$$
\mathrm{x} B 1 \mid \vdash \lambda \quad \text { or } \quad \mathrm{K} B 1 \mid \vdash-\lambda_{0}
$$

Suppose we consider telling it that it does not have all the teachers and get

```
KB2 = KB1 + \lambda.
```

If we now want to tell the $K B$ that Jim is a teacher, we get

[^1]The problem here is that KB3 still contains $\lambda$ end, consequently, still thinks it is missing a teacher. In fact, no matter how many teachers we tell the KB about, it will still think it is missing at least one. Moreover, if we try to tell it that it finally has all of them by adding $\boldsymbol{\lambda}$, then we arrive at an Inconsistent KB since it also contains $\lambda$.

What should have happened here is somewhat different. Once we arrived at KB2, the KB should know that it is missing a teacher and hence believe入. However, once a new teacher is added to the KB to produce KB3, the KB has no way of knowing whether or not this is the last teacher it was missing and so it should be the case (as with KB1) that neither

## RB3 if $\lambda$ nor KB3 |f $\rightarrow$.

In other words, after the introduction of Jim, the KB should no longer know whether or not it has all the teachers. In fact, the knowledge it had as KB2 is lost when it becomes KB3.

This is a strange kind of non-monotonicity. The usual symptom of a non-monotonic logic is that the addition of a new axiom Invalidates a previous theorem. In our case, the addition of the new axiom Invalidated a previous axiom. The curious puzzle here is that there is no "belief revision" going on in the sense of a realisation that an axiom was incorrectly added to the KB. Similarly, there is no admission of the world having changed in the sense of someone becoming or ceasing to be a teacher. In fact, the only change that has taken place is a change in what is known about the world. But this is enough since $\boldsymbol{\lambda}$ does make reference to the current state of the KB. So without admitting that the world has changed or that some previous statement about the world needs revision, we can still maintain that the truth value of $\lambda$ can change by noting that the state of the KB has changed. Thus, $\lambda$ cannot be part of the KB since the state of knowledge it refers to disappears as Information is acquired.

The solution to the problem of the addition of $\lambda$ to a KB is, therefore, to treat the formula as ordinary world knowledge where the K operator is used to refer to the current state of the KB. This is only natural since a $K B$ la assumed to have complete knowledge of itself. Conaequently, any mention of what is known in a new piece of information must be "referential" and not "attributive". Thus, the addition to the KB must be understood by first resolving these references. In other words, the solution to the problem is not to prohibit additions like $\boldsymbol{\lambda}$, but rather to allow $\boldsymbol{\lambda}$ in the KB definition language but not the KB itself. For example, KB2 now becomes

## KB2 - KB1 $+\lambda=$ \{Teachar (john) (Ex) [Tancher (x) A $\rightarrow(x=$ john $)]$ ]

where we have replaced the open formula

## KfTeacher(x)]

by a first order formula that resolves this
reference with respect to KB1:
(x-john).
This produces the property that
K 32 |f $\lambda$
and neither KB3 11 A nor KB3 | (- -,A as desired.

The solution I have proposed above presupposes that there will always be a formula of FOL that can be used to resolve any reference to what is known. It Is worth noting that this cannot be done Independently of a KB In that there is no formula $\alpha$ of FOL such that
$\vdash \mathbf{K} \equiv \mathbf{K a}$.
Of course, all we really need is a formula for each KB and not a formula that works for every KB. Fortunately, If we assume the availability of an equality predicate and restrict ourselves to finite KB's, this can always be done.

Theorean Aseme KB in finite and $\alpha(x l, \ldots, x k) \in K L$, There is a fortula $/ \alpha / \in$ FOL ( $w 1$ th equality) such that $K B$ if ( $x \perp$ )... ( $x k$ ) $(K a \equiv K / a /]$.

## The formula required here can be defined by:

/a/ = a when a e FOL, $/=a /=a / a /$,
$/(\alpha>\rho) /=(/ \alpha / v / \rho /), \quad /(x) \alpha /=(x) / a /$ and $/ \mathrm{k} \alpha /=\operatorname{Resolve}[/ \alpha /]$.

```
RESOLVE[a] = If a hat no free variables
    then if \overline{KB}\vdasha then (x)(x-x) elge q(x)(x-x)
            /* Acoume x fa free in a and that
                cl,...,ck but not c are all the constants
        1n a or KB */
    elae [(x-cl ^ RESOLVE[a<x/cl>]) v...v
        (x-ck ^ RESOLVE[\alpha<x/ck>]) v
        (x;c1 A...^^ x/ck ^ RESOLVE[u<x/c>]<c/x>)]
```

The method of allowing all of KL to specify a $K B$ is thus to let

## $K B+a=K B \cup\{/ a /\}$ for $a n y a \in K L$.

Viewed more semantics Uy In terms of world descriptions, we have that
$M(\alpha \beta) \cap\{\mathbf{w} \mid \mathcal{V}(\alpha, w, M(X B))-T\}=M(K B \cup(/ \alpha /\})$.
Note that although the function RESOLVE is not recursive, It is strictly proof-theoretic. Moreover, it can be shown that

which defines a syntactic version of query evaluation that is exactly equivalent to the semantic one defined earlier. Of course, the proof theory is not axiomatic but this is to be expected given that

## 

is not recursively enumerable. So, to summarise,
although the interaction with a first order KB should allow the language KL to be used, this Interaction can be understood in first order terms when the $K B$ is finite.

## 7. CONCLUSION

In this paper, I have considered from a formal standpoint the problem of interacting with an incomplete KB. The motivation behind the research is that, to effectively deal with partial knowledge, a system must first be able to determine the exact limits to what is known. This, In turn, places certain requirements on the language used to interact with a KB. In particular, the language has to be able to refer to the state of the $K B$ as well as to the state of the application domain.

To this effect, I proposed a language KL for interacting with a first order KB and presented a proof theory and semantics which were direct extensions of their FOL counterparts. KL was then applied to query evaluation and I showed that this required a non-monotonic operator. Semantic and syntactic interpretations of this operator were provided. Finally, KL was applied to KB definition and I demonstrated why this required reducing KL into the language of the KB, FOL. In fact, the net result was that the KB remained first order, but that interaction with it took place In a more expressive language.

Apart from the more practical implications of this work, there remain open questions even within the formal framework. I have not mentioned, for example, what impact the presence of an equality relation or non-constant terms would have on KL. Also, a method of handling defaults (and exceptions) is a reasonable goal given that the framework allows one to determine where their application is needed. In summary, the framework provides not only a formal standard against which to measure representation languages, but also a basis for further exploration.
[6] McCarthy, J., Sato, M., Hayashi, T. and Igarashl, S., "On the Model Theory of Knowledge", Memo AIM-312, Dept. of Computer Science, Stanford U., 1978.
[7] Vassiliou, Y., "A Formal Treatment of Imperfect Information in Database Management", Ph.D. thesis, Dept. of Computer Science, U. of Toronto, 1980.
[8] LIp8kl, V., "On Semantic Issues Connected with Incomplete Information Data Bases", Institute of Computer Science PAS Report 325, Warsaw, Poland, 1978.
[9] McDermott, D. and Doyle, J., "Non-monotonic Logic", Al memo 486, AI Lab., MIT, 1978.
[10] McDermott, D., "Non-monotonic Logic II: Nonmonotonic Modal Theories", Research Rep. No. 174, Dept. of Computer Science, Yale U., 1980.
[11] McCarthy, J., "Circumscription - A Form of Non-monotonic Reasoning", Artificial Intelligence 13, 1980.
[12] Reiter, R., "A Logic for Default Reasoning", Artificial Intelligence 13. 1980.
[13] Gallaire, H. and Minker, J. (eds.), Logic and Data Bases. Plenum Press, New York, 1978.
[14] Reiter, R., "On Closed World Data Bases" Logicc aind Dather Brases, Cbadlliárie e, H. and Minker, J. ((eddss).,), Plèeunra ripress, 1978

## ACKNOMFDGEMENTS

I would like to thank John Mylopoulos and Alex Borglda for contributing to the development of this work. I am also grateful to Teresa Miao for typing this manuscript.

## REFERENCES

[1] Reiter, R., "On Reasoning by Default", Proc. second TINLAP conference, Urbana, Illinois, 1978.
[2] Collins, A., Warnock, E., Alello, N. and Miller, M., "Reasoning from Incomplete Knowledge", Representation and Understanding: Studies in Cognitive Science, Collins, A. and Bobrow, D. (eds.), Academic Press, 1975.
[3] Levesque, H., "A Formal Treatment of Incomplete Knowledge Bases", Ph.D. thesis, Dept. of Computer Science, U. of Toronto, 1981.
[4] Hintikka, J., Knowledge and Belief; An Introduction to the Logic of the Two Notions. Cornell University Press, 1962.
[5] Moore, R., "Reasoning about Knowledge and Action", Ph.D. thesis, Dept. of E.E. and Computer Science, MIT, 1980.


[^0]:    This work was sponsored in part by the Natural Sciences and Engineering Council of Canada and the Department of Computer Science, U. of Toronto.

[^1]:    KB3 - KB2 + Teacher (Jim).

