



Centrum voor Wiskunde en Informatica
Centre for Mathematics and Computer Science

M. Bezem

Semantics and consistency of rule-based expert systems

Computer Science/Department of Software Technology

Report CS-R8824

July

The Centre for Mathematics and Computer Science is a research institute of the Stichting Mathematisch Centrum, which was founded on February 11, 1946, as a nonprofit institution aiming at the promotion of mathematics, computer science, and their applications. It is sponsored by the Dutch Government through the Netherlands Organization for the Advancement of Pure Research (Z.W.O.).

Copyright © Stichting Mathematisch Centrum, Amsterdam

69 K 13, 69 K 14, 69 F 41

Semantics and Consistency of Rule-based Expert Systems

Marc Bezem

*Centre for Mathematics and Computer Science
P.O. Box 4079, 1009 AB Amsterdam, The Netherlands*

Consistency of a knowledge-based system has become a topic of growing concern. Every notion of consistency presupposes a notion of semantics. We present a theoretical framework in which both the semantics and the consistency of a knowledge base can be studied. This framework is based on first order flat many-sorted predicate logic and is sufficiently rich to capture an interesting class of rule-based expert systems and deductive databases. We analyse the feasibility of the consistency test and prove that this test is feasible for knowledge bases in Horn format without quantification. This report is a revised and extended version of report CS-R8736. The extension concerns the use of expressions in rules.

1980 Mathematics Subject Classification: 68T30, 68T15.

1987 CR Categories: I.2.3, I.2.4, F.4.1.

Key Words & Phrases: knowledge-based systems, rule-based expert systems, knowledge representation, consistency.

Note: The material for this article has been taken from the PRISMA documents P176 and P297. The research is part of the PRISMA project (PaRallel Inference and Storage MAchine), a joint effort with Philips Research Eindhoven, partially supported by the Dutch "Stimulerings-projectteam Informatica-onderzoek" (SPIN).

INTRODUCTION

The plan of this paper (a revised and extended version of [B]) is as follows. First the reader is introduced to the knowledge representation used in rule-based expert systems. We shall indicate some semantical problems in relation to this knowledge representation. Then we explain in an informal way how many-sorted predicate logic comes in. In the next section we describe syntax and semantics of many-sorted predicate logic. We assume some knowledge of first order logic. Thereafter we shall be able to characterize rule-based expert systems as first order theories. The Tarski semantics solves the semantical problems mentioned above. Furthermore we shall derive several results on decidability and consistency of rule-based expert systems. Unfortunately, some natural equality and ordering axioms are not in Horn format (see [Re] for a discussion on the domain closure axiom). Hence testing consistency with a standard theorem prover would be very inefficient. In the fourth section we describe a technical device, a certain kind of null value, which allows feasible consistency testing in the presence of equality and ordering axioms, which are not in Horn format. The underlying idea is partiality of functions, thus avoiding the search for consistent function values in a (possibly gigantic) product space. This kind of null value, being quite different from null values as described in [IL], appears to be new. In the last section we show how to extend our results to knowledge bases with arithmetical (and other) expressions in the conditions of the rules. We shall focus our attention on rule-based expert systems, but the techniques can also be applied to deductive databases with functional dependencies.

Report CS-R8824
Centre for Mathematics and Computer Science
P.O. Box 4079, 1009 AB Amsterdam, The Netherlands

1. RULE-BASED EXPERT SYSTEMS

1.1. In rule-based expert systems shells such as EMYCIN [BS] or DELFI2 [L], knowledge about some specific domain can be expressed in *facts* and in *rules* of the form

IF antecedent THEN consequent.

Facts are so-called object-attribute-value triples, or $\langle o, a, v \rangle$ triples for short. The antecedent of a rule is a conjunction of disjunctions of conditions, and conditions are definite statements, such as *same*, *notsame* and *less than*, about $\langle o, a, v \rangle$ triples. We restrict ourselves to rules having as consequent a conjunction of conclusions of the form *conclude* $\langle o, a, v \rangle$. In most cases so-called certainty factors are associated with the facts and the rules. Certainty factors range from 1.00 (definitely true) to -1.00 (definitely false). The certainty factor of a fact expresses a measure of certainty about that fact, whereas the certainty factor of a rule scales the measure of certainty about the consequent with respect to the measure of certainty about the antecedent. In DELFI2 an *object tree* (called *context tree* in MYCIN) is used to state properties of and relations between different objects, which cannot be expressed by the rules. The nodes of this tree are objects, labeled by their attributes and respective domains of values. The path from a node to the root of this tree constitutes the context of that node. In other words: the objects occurring in the subtree of a node are sub-objects of the object belonging to that node. Furthermore it is stated in the object tree whether an attribute is singlevalued or multivalued.

The interpretation of the knowledge in rule-based expert systems is more operational than declarative: *same* $\langle o, a, v \rangle$ is true if and only if $\langle o, a, v \rangle$ occurs as fact (with certainty factor > 0.2), *conclude* $\langle o, a, v \rangle$ has the effect that $\langle o, a, v \rangle$ is added as fact (with appropriate certainty factor), and if the antecedent of a rule evaluates to true, then that rule may be *fired*, i.e. all conclusions occurring in the consequent are executed. We remark that *same* $\langle o, a, v \rangle$ and *conclude* $\langle o, a, v \rangle$ have the same declarative meaning as the fact $\langle o, a, v \rangle$, i.e. attribute *a* of object *o* has value *v*.

1.2. Consider the following real-life example extracted from HEPAR, an expert system for the diagnosis of liver and biliary disease, built with DELFI2.

IF same $\langle \textit{patient}, \textit{complaint}, \textit{colicky_pain} \rangle$
THEN conclude $\langle \textit{patient}, \textit{pain}, \textit{colicky} \rangle$ (1.00)

IF same $\langle \textit{patient}, \textit{abd_pain}, \textit{yes} \rangle$ AND
same $\langle \textit{pain}, \textit{character}, \textit{continuous} \rangle$
THEN conclude $\langle \textit{patient}, \textit{pain}, \textit{colicky} \rangle$ (-1.00)

IF same $\langle \textit{patient}, \textit{complaint}, \textit{abdominal_pain} \rangle$ OR
same $\langle \textit{patient}, \textit{pain}, \textit{colicky} \rangle$
THEN conclude $\langle \textit{patient}, \textit{abd_pain}, \textit{yes} \rangle$ (1.00)

These three rules (from a rule base consisting of over 400 rules) show two *objects*, *patient* and *pain*, four *attributes*, namely *complaint*, *pain* and *abd_pain* of *patient* and *character* of *pain*, as well as several *values*.

1.3. A first observation which can be made is the following. Five items in the rules above refer to *pain*: the values *colicky_pain* and *abdominal_pain*, the attributes *pain* and *abd_pain* of the object *patient*, and the object *pain*, which is a sub-object of *patient*, as stated by the object tree of HEPAR.

The interrelations between these items do not seem to be expressible by the formalism.

A second observation on the three rules above is their inconsistency in the presence of the facts $\langle \textit{patient}, \textit{complaint}, \textit{colicky_pain} \rangle$ (1.00) and $\langle \textit{pain}, \textit{character}, \textit{continuous} \rangle$ (1.00). The inference engine reasons backwards, using a simple loop check to prevent infinite looping. Depending on the presence of the fact $\langle \textit{patient}, \textit{complaint}, \textit{abdominal_pain} \rangle$ (1.00), the inference engine did or did not hit upon the contradiction $\langle \textit{patient}, \textit{pain}, \textit{colicky} \rangle$ (1.00 and -1.00). In *both* cases the contradiction was completely ignored.

These observations show a defect of the knowledge representation used in rule-based expert systems, namely the lack of a clear semantics.

1.4. Basically our approach amounts to interpreting $\langle o, a, v \rangle$ by $a(o, v)$ in the multivalued, and by $a(o) = v$ in the singlevalued case. Here o and v are constants for elements of a domain O of objects and a domain V of values. In the multivalued case a denotes a *relation*, i.e. a subset of $O \times V$, and in the singlevalued case a *function* from O to V . If o is a sub-object of o' , then o' is added as argument of a .

1.5. The following examples show how to extend our interpretation of $\langle o, a, v \rangle$ triples to atoms.

conclude $\langle \textit{patient}, \textit{complaint}, \textit{abdominal_pain} \rangle$

becomes

complaint ($\textit{patient}, \textit{abdominal_pain}$)

less_than $\langle \textit{patient}, \textit{temperature}, 36.8 \rangle$

becomes

temperature ($\textit{patient}$) < 36.8

same $\langle \textit{pain}, \textit{character}, \textit{continuous} \rangle$

becomes

character ($\textit{patient}, \textit{pain}$) = *continuous*

Note that the fact that pain is a sub-object of patient is expressed by adding *patient* as an argument of the function *character*. Another example, based on the rules of 1.2, is to be found in 4.6.

We feel that the existing formalisms for handling uncertainty are unsatisfactory. As we do not have a good alternative we leave the subject aside.

1.6. Under the interpretation described above, a rule-based expert system becomes a theory in first order many-sorted predicate logic, in short: a many-sorted theory. To keep this paper self-contained we give a short, introductory description of the syntax and semantics of many-sorted predicate logic in the next section. In Section 3 we shall characterize expert systems as many-sorted theories of a certain type. This approach has the following advantages:

- The declarative semantics of the expert system becomes perfectly clear, being the Tarski semantics of the associated many-sorted theory.
- Logical concepts such as consequence, consistency and decidability get a clear meaning in relation to the expert system.
- Theorem proving techniques for testing consistency, such as resolution, become available for the expert system.

The choice for many-sorted instead of one-sorted logic is motivated as follows:

- It is natural to make a distinction between numerical data (ordered by $<$) and symbolic data (usually unordered).
- Subdividing a set of constants into sorts and typing the predicate and function symbols considerably reduces the number of well-formed formulas, which is important for debugging large knowledge bases.
- In the presence of variables the many-sorted approach imposes restrictions on unification, which considerably reduces the resolution search space (see for example [W]). We shall not make use of this advantage in the present paper.

2. MANY-SORTED PREDICATE LOGIC

2.1. The syntax of many-sorted predicate logic extends the syntax of ordinary, one-sorted, predicate logic by having a set of *sorts* Σ , instead of just one sort. Moreover we have *variables* x_i^σ and *constants* c_i^σ for all sorts $\sigma \in \Sigma$. Furthermore we have *function symbols* $f_i^{\sigma_1 \times \dots \times \sigma_m \rightarrow \sigma_0}$ ($m > 0$), where the notion of *type* $\sigma_1 \times \dots \times \sigma_m \rightarrow \sigma_0$ replaces the notion of *arity* from the one-sorted case. We also have *proposition symbols* p_i and *predicate symbols* $P_i^{\sigma_1 \times \dots \times \sigma_m}$ ($m > 0$) of type $\sigma_1 \times \dots \times \sigma_m$. *Terms* are formed from variables and constants by function application (respecting the sorts). *Atoms* are either proposition symbols or the application of a predicate symbol to terms of appropriate sorts. With the help of propositional connectives and quantifiers, atoms are combined into *formulas*. The sets Σ , CONS, FUNC, PROP and PRED of, respectively, sorts, constants, function symbols, proposition symbols and predicate symbols, form together the *similarity type* of a specific many-sorted predicate calculus. For practical reasons we assume that the similarity type is finite.

2.2. A *many-sorted structure* \mathfrak{N} consists of:

- (a) A non-empty set A_σ for each $\sigma \in \Sigma$, called the *domain* of sort σ .
- (b) For each constant c_i^σ an element $\bar{c}_i \in A_\sigma$.
- (c) For each function symbol $f_i^{\sigma_1 \times \dots \times \sigma_m \rightarrow \sigma_0}$ a mapping $\bar{f}_i : A_{\sigma_1} \times \dots \times A_{\sigma_m} \rightarrow A_{\sigma_0}$.
- (d) For each proposition symbol p_i a truth value \bar{p}_i .
- (e) For each predicate symbol $P_i^{\sigma_1 \times \dots \times \sigma_m}$ a mapping $\bar{P}_i : A_{\sigma_1} \times \dots \times A_{\sigma_m} \rightarrow \{TRUE, FALSE\}$.

2.3. An *assignment* in \mathfrak{N} is a mapping a assigning to each variable x_i^σ an element $a(x_i^\sigma)$ of A_σ .

2.4. The *interpretation* in \mathfrak{N} of a term t under an assignment a , denoted by $I_a^{\mathfrak{N}}(t)$ or \bar{t} for short, is inductively defined as follows:

- (a) $\bar{x}_i^\sigma = a(x_i^\sigma)$
- (b) $\bar{c}_i^\sigma = \bar{c}_i$
- (c) $\bar{f}_i^{\sigma_1 \times \dots \times \sigma_m \rightarrow \sigma_0}(t_1, \dots, t_m) = \bar{f}_i(\bar{t}_1, \dots, \bar{t}_m)$

The *truth value* in \mathfrak{N} of an atom $P_i^{\sigma_1 \times \dots \times \sigma_m}(t_1, \dots, t_m)$ under an assignment a is given by $\bar{P}_i(\bar{t}_1, \dots, \bar{t}_m)$.

2.5. The *truth value* in \mathfrak{N} of a formula F under an assignment a , denoted by $\mathcal{V}_a^{\mathfrak{N}}(F)$, is inductively defined as follows:

- (a) If F is an atom, then $\mathcal{V}_a^{\mathfrak{N}}(F)$ is given by 2.4.
- (b) $\mathcal{V}_a^{\mathfrak{N}}$ respects the truth tables of the propositional connectives.
- (c) $\mathcal{V}_a^{\mathfrak{N}}(\forall x_i^\sigma F) = TRUE$ if and only if for all assignments a' , which differ at most on x_i^σ from a , we have $\mathcal{V}_{a'}^{\mathfrak{N}}(F) = TRUE$.
- (d) $\mathcal{V}_a^{\mathfrak{N}}(\exists x_i^\sigma F) = TRUE$ if and only if there exists an assignment a' , which differs at most on x_i^σ from a , such that $\mathcal{V}_{a'}^{\mathfrak{N}}(F) = TRUE$.

2.6. A formula F is *true* in \mathfrak{M} , denoted by $\models_{\mathfrak{M}} F$, if $\forall_a^{\mathfrak{M}}(F) = TRUE$ for all assignments a .

2.7. A *sentence* (or *closed* formula) is a formula without free variables (i.e. variables which are not bound by a quantifier). It will be clear that for sentences S the truth value $\forall_a^{\mathfrak{M}}(S)$ does not depend on the assignment a . As a consequence we have either $\models_{\mathfrak{M}} S$ or $\models_{\mathfrak{M}} \neg S$, for every sentence S . Let SENT denote the set of sentences.

2.8. Let $\Gamma \subset \text{SENT}$. \mathfrak{M} is called a *model* for Γ , denoted by $\models_{\mathfrak{M}} \Gamma$, if $\models_{\mathfrak{M}} S$ for all $S \in \Gamma$.

2.9. $S \in \text{SENT}$ is called a (*semantical*) *consequence* of (or *implied* by) $\Gamma \subset \text{SENT}$ if for all many-sorted structures \mathfrak{M} we have: if $\models_{\mathfrak{M}} \Gamma$, then $\models_{\mathfrak{M}} S$. This will be denoted by $\Gamma \vDash S$ (or $\vDash S$ if Γ is empty). Furthermore we define the *theory* of Γ as the set $Th(\Gamma) = \{S \in \text{SENT} \mid \Gamma \vDash S\}$.

2.10. $\Gamma \subset \text{SENT}$ is called *consistent* if Γ has a model. $Th(\Gamma)$ is called *decidable* if there exists a mechanical decision procedure to decide whether a given sentence S is a semantical consequence of Γ or not.

2.11. Two many-sorted structures are called *elementarily equivalent* if exactly the same sentences are true in both structures.

2.12. REMARKS.

2.12.1. We refrain from giving an axiomatization of many-sorted predicate logic since our main concern will be model theory. Most textbooks on mathematical logic provide a complete axiomatization of ordinary (one-sorted) predicate logic. It suffices to generalize the quantifier rules in order to obtain an axiomatization of flat many-sorted predicate logic.

2.12.2. Of course, one-sorted predicate logic is a special case of many-sorted predicate logic. As a consequence, the latter is as undecidable as the former. More precisely: $\vDash S$ is undecidable, provided that the similarity type is rich enough (CHURCH, TURING, 1936, see also [M, 16.58]).

2.12.3. Conversely, many-sorted predicate logic can be embedded in one-sorted predicate logic by adding unary predicate symbols $S^\sigma(x)$, expressing that x is of sort σ , and replacing inductively in formulas $\forall x_i^q F$ (resp. $\exists x_i^q F$) by $\forall x(S^\sigma(x) \rightarrow F)$ (resp. $\exists x(S^\sigma(x) \wedge F)$). Let A' be the one-sorted sentence obtained from $A \in \text{SENT}$ in this way. It can be proved (see [M]) that $\vDash A$ if and only if $\Gamma \vDash A'$, where $\Gamma = \{\exists x S^\sigma(x) \mid \sigma \in \Sigma\}$ expresses the fact that the domains are non-empty. This embedding allows us to generalize immediately many results on one-sorted predicate calculus to the many-sorted case (e.g. the compactness theorem). We shall not make use of this possibility in the present paper.

3. RULE-BASED EXPERT SYSTEMS AS MANY-SORTED THEORIES

3.0. We propose the following terminology for certain kinds of many-sorted theories:

- Indexed propositional expert systems.
- Universally quantified expert systems.

3.1. An *indexed propositional expert system* is a many-sorted theory axiomatized by:

- (a) *Explicit axioms* (the rule base and the fact base), which are boolean combinations of atoms of the form $P^{\sigma_1 \times \dots \times \sigma_n}(c, \dots, c')$ or of the form $f^{(\sigma_1 \times \dots \times \sigma_n) \rightarrow \sigma_0}(c, \dots, c') =_{\sigma_0} c''$ (resp. $<_{\sigma_0} c'', >_{\sigma_0} c''$), with constants c, \dots, c', c'' of appropriate sorts. Such atoms (here and below called *constant-atoms*, or *c-atoms* for short) may be viewed as indexed propositions, which explains the name. Note that we conform to the convention to write $=$, $<$ and $>$ as infix predicates.

- (b) *Implicit axioms* for equality of each sort and ordering of each sort for which an ordering is appropriate. The axioms for $=_\sigma$, equality of sort σ , are (loosely omitting sort super- and subscripts):

$$\begin{aligned} & \forall x(x = x), \\ & \forall x_1, x_2(x_1 = x_2 \rightarrow x_2 = x_1), \\ & \forall x_1, x_2, x_3((x_1 = x_2 \wedge x_2 = x_3) \rightarrow x_1 = x_3), \\ & \forall x_1, x_2((x_1 = x_2 \wedge F(x_1)) \rightarrow F(x_2)) \text{ for formulas } F, \\ & \neg c_i = c_j \text{ for } 0 \leq i \neq j \leq n, \\ & \forall x(x = c_0 \vee \dots \vee x = c_n), \text{ the domain closure axiom.} \end{aligned}$$

These axioms express that $=$ is a congruence relation on a finite domain, where every element has exactly one name. Let EQ denote the set of equality axioms for all sorts σ , then we have by definition that either $EQ \models c_i = c_j$ or $EQ \models \neg c_i = c_j$ for all i, j . The axioms for $<$ and $>$ are:

$$\begin{aligned} & \forall x_1, x_2, x_3((x_1 < x_2 \wedge x_2 < x_3) \rightarrow x_1 < x_3), \\ & \forall x(\neg x < x), \\ & \forall x_1, x_2(x_1 < x_2 \vee x_1 = x_2 \vee x_2 < x_1) \text{ for sorts which are totally ordered,} \\ & \forall x_1, x_2(x_1 < x_2 \leftrightarrow x_2 > x_1), \\ & \text{a subset } \Delta \text{ of } \{c_i < c_j \mid 0 \leq i \neq j \leq n\} \cup \{\neg c_i < c_j \mid 0 \leq i, j \leq n\} \text{ (see below).} \end{aligned}$$

These axioms express that $<$ is a transitive, irreflexive and (possibly) total ordering with inverse $>$. Let O denote the set of ordering axioms. We require that Δ is such that either $O \models c_i < c_j$ or $O \models \neg c_i < c_j$, for all i, j . It follows in particular that $EQ \cup O$ is consistent.

The idea behind the implicit axioms is that $=$ and $<$ are provided by the system and have a fixed meaning, whereas the other predicates are user-defined. We put $IA = EQ \cup O$, so IA denotes the set of implicit axioms.

3.2. A *universally quantified expert system* differs from an indexed propositional one by allowing not only constants, but also variables in the explicit axioms. All explicit axioms are assumed to be universally closed.

3.3. Among the theories that do not fall under 3.1 and 3.2 are many theorem provers. A theorem prover might be an undecidable theory. The theoretical observations below show that, from a logical point of view, expert systems as defined in 3.1 and 3.2 are very simple, decidable theories.

3.4. LEMMA. *For every $S \in \text{SENT}$ there exists a boolean combination S' of closed atoms such that $EQ \models S \leftrightarrow S'$.*

PROOF. Replace inductively every subformula $\forall x F(x)$ by $F(c_0) \wedge \dots \wedge F(c_n)$ and $\exists x F(x)$ by $F(c_0) \vee \dots \vee F(c_n)$ where $\forall x(x = c_0 \vee \dots \vee x = c_n) \in EQ$. \square

3.5. LEMMA. *For every closed atom A there exists a boolean combination A' of c -atoms such that $EQ \models A \leftrightarrow A'$.*

PROOF by giving a typical example. Let A be for instance $P(f(f'(c^{\sigma_0}), c^{\sigma_1}), c^{\sigma_2})$ with f of type $\sigma_2 \times \sigma_1 \rightarrow \sigma_0$ and f' of type $\sigma_0 \rightarrow \sigma_2$. Let A^\exists be the sentence $\exists x^{\sigma_2} \exists y^{\sigma_0} [f'(c^{\sigma_0}) = x^{\sigma_2} \wedge f(x^{\sigma_2}, c^{\sigma_1}) = y^{\sigma_0} \wedge P(y^{\sigma_0}, c^{\sigma_2})]$. Now apply Lemma 3.4 to A^\exists and obtain A' . \square

3.6. Both lemmas above can cause combinatorial explosions. Therefore they are only of theoretical use. They tell us that, in the presence of EQ , boolean combinations of c -atoms have the same expressive power as full many-sorted predicate logic.

3.7. LEMMA. *If $EQ \subset \Gamma \subset \text{SENT}$, then every model of Γ is elementarily equivalent to a model whose domains consist of exactly the interpretations of all constants.*

PROOF. Let \mathfrak{M} be a model of Γ with $EQ \subset \Gamma \subset \text{SENT}$. By EQ the interpretations of the equality predicates in \mathfrak{M} are equivalence relations which are congruences with respect to the interpretation of all other predicate and function symbols in \mathfrak{M} . It follows that \mathfrak{M} and $\mathfrak{M}/=$, the quotient structure of \mathfrak{M} modulo equality, are elementarily equivalent. By the axioms $\forall x(x=c_0 \vee \dots \vee x=c_n)$ and $\neg c_i=c_j$ ($i \neq j$) from EQ , it follows that the domains of $\mathfrak{M}/=$ consist of exactly the interpretations of all constants. \square

3.8. LEMMA. *If $EQ \subset \Gamma \subset \text{SENT}$, then $\text{Th}(\Gamma)$ is decidable.*

PROOF. Let Γ be such that $EQ \subset \Gamma \subset \text{SENT}$. Then we can apply Lemma 3.7 and observe that there are just finitely many non-isomorphic $\mathfrak{M}/=$'s. In other words: up to dividing out $=$ and identifying individuals with the same name there are just finitely many different models of Γ . Moreover all models of Γ are finite. Hence we can, for any given $S \in \text{SENT}$, test in finite time whether S holds in all models of Γ or not. \square

4. TESTING CONSISTENCY

4.1. Definition 2.10 and Lemma 3.7 suggest the following procedure for testing the consistency of theories Γ with $EQ \subset \Gamma$: generate all many-sorted structures whose domains consist of exactly the interpretations of all constants, and test each time whether the many-sorted structure is a model of Γ or not. This procedure is in general not feasible. A first step towards a feasible consistency test is the alternative characterization of consistency for indexed propositional expert systems, described in the following paragraphs.

Let Γ be an axiomatization of an indexed propositional expert system (see 3.1). Let P_1, \dots, P_m be a list of all c -atoms of the form $P(c, \dots, c')$ occurring in Γ , and let t_1, \dots, t_n list all terms $f(c, \dots, c')$ occurring in Γ . The idea behind the following construction is that it is not necessary to have an entire many-sorted structure to be able to interpret Γ .

Let $A_\sigma = \{c_0^\sigma, \dots, c_n^\sigma\}$ be the set of all constants of sort $\sigma \in \Sigma$. Equality of sort σ is interpreted by syntactical identity on A_σ . Then all equality axioms from EQ are satisfied. The ordering (if any) on A_σ is induced by O , i.e. $c_i < c_j$ if and only if $O \vDash c_i < c_j$. We need the following notions:

- A *truth valuation* of P_1, \dots, P_m is an assignment of either *TRUE* or *FALSE* to each P_i ($1 \leq i \leq m$).
- A *valuation* of t_1, \dots, t_n is an assignment of a unique element of the appropriate domain A_σ to each t_j ($1 \leq j \leq n$). It will be clear that a truth valuation of P_1, \dots, P_m and a valuation of t_1, \dots, t_n suffice for an interpretation of Γ . Each model of Γ yields a truth valuation of P_1, \dots, P_m and a valuation of t_1, \dots, t_n . Conversely, each truth valuation of P_1, \dots, P_m and valuation of t_1, \dots, t_n for which every explicit axiom of Γ is true, can be extended in an arbitrary way to a many-sorted structure $\mathfrak{M} \vDash \Gamma$. Thus we have the following

THEOREM. *Let conditions be as above. Then we have: Γ is consistent if and only if there exists a truth valuation of P_1, \dots, P_m and a valuation of t_1, \dots, t_n for which every explicit axiom of Γ is true.*

4.2. Theorem 4.1 suggests a simple algorithm for testing the consistency of an indexed propositional expert system Γ : generate all valuations and truth valuations and test each time whether the explicit

axioms of Γ are satisfied or not. This clearly leads to combinatorial explosions. A second step towards a feasible consistency test is imposing language restrictions on our expert systems.

A common language restriction is Horn format. A *clause* is a finite disjunction of atoms and negations of atoms (so called positive and negative *literals*). A *conjunctive normal form* is a finite conjunction of clauses. A *Horn clause* is a clause containing at most one positive literal. A *unit clause* is a clause consisting of one positive literal. The connection between Horn clauses and production rules is easily seen by the equivalence of $(A_1 \wedge \dots \wedge A_k) \rightarrow B$ and $A_1^- \vee \dots \vee A_k^- \vee B^+$, where the superscripts $+$ and $-$ denote whether a literal occurs positively or negatively. However, the implicit axioms $\forall x(x = c_0 \vee \dots \vee x = c_n)$ and $\forall x_1, x_2(x_1 < x_2 \vee x_1 = x_2 \vee x_2 < x_1)$ are not Horn clauses. As a consequence we can only require the explicit axioms to be Horn clauses. This is not sufficient for feasible consistency checking, as will be demonstrated in the next paragraph. The idea is to reduce the satisfiability problem for propositional logic, which is known to be NP-complete (see [GJ]), to the consistency problem of indexed propositional expert systems, all whose explicit axioms are Horn clauses.

Let σ be a sort with exactly two constants: c_0^σ and c_1^σ . Then we have $\forall x(x = c_0 \vee x = c_1) \in EQ$. For any propositional atom A , let f_A be a function symbol of type $\sigma \rightarrow \sigma$. It is not difficult to see that

$$\{\neg f_A(c_0) = c_0 \vee A, f_A(c_0) = c_0 \vee \neg A\} \cup EQ \models (\neg f_A(c_0) = c_1) \leftrightarrow A.$$

So the *positive* literal A is equivalent to the *negative* literal $\neg f_A(c_0) = c_1$ in any Γ containing the two Horn clauses $\neg f_A(c_0) = c_0 \vee A$ and $f_A(c_0) = c_0 \vee \neg A$ and the equality axioms EQ . Let C be any propositional conjunctive normal form. Let Γ_C be the indexed propositional expert system with explicit axioms $\neg f_A(c_0) = c_0 \vee A$ and $f_A(c_0) = c_0 \vee \neg A$ for all positive literals A occurring in C . Then Γ_C is consistent and satisfies $\Gamma_C \models C \leftrightarrow H_C$, where H_C is the set (conjunction) of Horn clauses obtained from C by replacing all positive literals A by their equivalent negative literal. Moreover we have that C is satisfiable if and only if $\Gamma_C \cup H_C$ is consistent. This yields a polynomial reduction of the satisfiability problem for propositional logic to the consistency problem of indexed propositional expert systems, all whose explicit axioms are Horn clauses. A similar reduction could be established using $\forall x_1, x_2(x_1 < x_2 \vee x_1 = x_2 \vee x_2 < x_1)$ instead of $\forall x(x = c_0 \vee x = c_1)$.

4.3. In the previous subsection we showed that testing consistency is NP-hard without further language restrictions. The problem is to specify language restrictions, which are strong enough to guarantee a feasible consistency test, and still allow enough expressivity for some given application. This third and final step towards a feasible consistency test will be achieved in the following theorem by the introduction of constants *undefined*, which are different from (wrt. $=$) and incomparable with (wrt. $<$) any other constant. These constants are used for the simulation of partiality of functions, thus avoiding the search for consistent function values in a (possibly gigantic) product space of sorts. Further intuition will be given in 5.4.1.

THEOREM. *Let Γ be the axiomatization of an indexed propositional expert system satisfying the following conditions:*

- (1) *All explicit axioms of Γ are Horn clauses.*
- (2) *For every sort $\sigma \in \Sigma$ there exists at least one constant, which does not occur in the explicit axioms of Γ . Such a constant, say c_0^σ , will be denoted by *undefined* $^\sigma$. So we have:
 $\forall x(x = \text{undefined} \vee x = c_1 \vee \dots \vee x = c_n) \in EQ$ and $\neg \text{undefined} = c_i \in EQ$ for all $1 \leq i \leq n$.*
- (3) *For every sort $\sigma \in \Sigma$ the ordering (if any) of sort σ is partial with respect to *undefined* $^\sigma$. More precisely: we require that $O \models \neg \text{undefined} < c_i$ and $O \models \neg \text{undefined} > c_i$ for all $1 \leq i \leq n$.*
- (4) *The orderings $<$ and $>$ do not occur in the positive literals occurring in the explicit axioms of Γ .*

Then the consistency of Γ can be tested in polynomial time.

PROOF. List the explicit axioms of Γ as follows (the super- and subscripted capitals denote c -atoms):

$$\begin{aligned}
& A_1^+ \\
& \cdot \\
& \cdot \\
& A_p^+ \\
& C_{1,1}^- \vee \cdots \vee C_{1,m_1}^- \vee D_1^+ \\
& \cdot \\
& \cdot \\
& C_{q,1}^- \vee \cdots \vee C_{q,n_q}^- \vee D_q^+ \\
& B_{1,1}^- \vee \cdots \vee B_{1,m_1}^- \\
& \cdot \\
& \cdot \\
& B_{r,1}^- \vee \cdots \vee B_{r,m_r}^-
\end{aligned}$$

Now apply the following algorithm, which is in fact a special case of hyper-resolution (see [R]). Under the conditions of the theorem, positive literals are either of the form $P(c, \dots, c')$, or of the form $f(c, \dots, c') = c''$. The current set of unit clauses is denoted by U_s . Anticipating the generalization in Section 5 the formulation of the algorithm is slightly more general than necessary: " $EQ \cup U_s$ is consistent" simply says that we do not have $f(c, \dots, c') = c_i \in U_s$ and $f(c, \dots, c') = c_j \in U_s$ with $i \neq j$, and " c -atom X is implied by $IA \cup U_s$ " is the case if either $X \in U_s$, or X is of the form $f(c, \dots, c') < c_i$ (resp. $> c_i$) and we have $f(c, \dots, c') = c_j \in U_s$ with $0 \vDash c_j < c_i$ (resp. $0 \vDash c_j > c_i$).

(* generation of unit clauses *)

$$U_s := \{A_1^+, \dots, A_p^+\};$$

$$U_{new} := \{\};$$

REPEAT

$$U_s := U_s \cup U_{new};$$

$$U_{new} := \{\};$$

IF $EQ \cup U_s$ is consistent

THEN

FOR each clause $C_{i,1}^- \vee \cdots \vee C_{i,n_i}^- \vee D_i^+$

DO

cancel all $C_{i,j}^-$ such that $C_{i,j}^+$ is implied by $IA \cup U_s$;

IF all negative literals of the clause are cancelled

THEN $U_{new} := U_{new} \cup \{D_i^+\}$

OD

UNTIL

$$U_{new} \subset U_s;$$

(* test if the unit clauses form a model *)

IF $EQ \cup U_s$ is consistent AND there exists no clause $B_{i,1}^- \vee \dots \vee B_{i,m}^-$ such

that all $B_{i,j}^+$ are implied by $IA \cup U_s$,

THEN Γ is consistent

ELSE Γ is inconsistent

The algorithm terminates since the number of literals is strictly decreasing. By the cancellations the explicit axioms of Γ are transformed into an equivalent set of Horn clauses. The consistency in the THEN-part can be seen as follows. First note that $EQ \cup U_s$ is consistent. Secondly, every clause not in U_s contains a negative literal L^- whose complement L^+ is not implied by $IA \cup U_s$. It can happen that $IA \cup U_s$ implies L^- (for example $f(c)=2 \in U_s$, $L^- = \neg f(c) < 1$, $1 < 2 \in \Delta \subset O$). Then every clause containing L^- holds in any model of $IA \cup U_s$. If neither L^- nor L^+ is implied by $IA \cup U_s$, then we have either $L^+ = P(c, \dots, c') \notin U_s$, or L^+ contains a function term $f(c, \dots, c')$ such that $f(c, \dots, c') = c'' \in U_s$ for no constant c'' . We define the following (truth-) valuation (see 4.1):

$$P_i = \begin{cases} TRUE & \text{if } P_i \in U_s \\ FALSE & \text{otherwise} \end{cases}$$

$$t_j = \begin{cases} c & \text{if } t_j = c \in U_s \\ \text{undefined} & \text{otherwise} \end{cases}$$

The valuation is well-defined since $EQ \cup U_s$ is consistent. By the properties of the constants *undefined* (conditions (2) and (3) of the theorem), the literals L^- mentioned above are *TRUE* in the (truth-) valuation. Hence every clause containing L^- is *TRUE* in the above (truth-) valuation, which constitutes (in the sense of Theorem 4.1) a model of the entire set of clauses. The algorithm is clearly quadratic in the number of occurrences of literals. To get the consistency of Γ in complexity class P we tacitly assumed that the primitive operations used in the algorithm above are in P. \square

REMARKS.

4.3.1. As follows by close inspection of the proof above, it would suffice to require the following weakening of condition (2): for every sort σ for which a function symbol f of type $\dots \rightarrow \sigma$ occurs in a negative literal occurring in the explicit axioms of Γ , there exists at least one constant which does not occur on the right-hand side of an equation occurring in such a literal. However, we think it is more systematic to require condition (2) as it stands. In practice the constants *undefined* are simply added to every domain of constants occurring in the knowledge base.

4.3.2. Condition (4) can not be missed, which can be seen as follows. Assume for some sort σ we have exactly three constants different from *undefined*, which are totally ordered by $c_1 < c_2 < c_3$. Then the unit clause $f(c) > c_1$ is equivalent to $f(c) = c_2 \vee f(c) = c_3$, which enables a similar construction as in the third paragraph of 4.2.

4.3.3. In view of condition (1), occurrences of *notsame* $\langle o, a, v \rangle$ in the antecedent of a rule can be problematic, since negative conditions in a rule result in positive literals in the clausal form. This problem can be overcome by postponing the consistency test until these occurrences evaluate to either *TRUE* or *FALSE* (e.g. after querying the user). Then the knowledge base can be transformed into an equivalent knowledge base satisfying (1). An alternative would be to allow clauses with more than one positive literal and to apply the following lemma. Of course the consistency test then becomes in general NP-hard, the worst-case time complexity doubles with every application of the lemma and we can only hope that practical cases are not the worst.

LEMMA. Let Γ be an axiomatization of an indexed propositional expert system. Let furthermore, for some c -atom A , Γ_1 be obtained from Γ by omitting (an arbitrary number of) explicit axioms which are in clausal form (not necessarily Horn format) and contain A as positive literal, and Γ_2 by deleting A from (an arbitrary number of) such axioms. Then we have: Γ is consistent if and only if $\{A\} \cup \Gamma_1$ is consistent or Γ_2 is consistent.

PROOF. As to the if-part we remark that $\{A\} \cup \Gamma_1 \models \Gamma$ and $\Gamma_2 \models \Gamma$. As to the only-if-part, note that for every many sorted structure \mathfrak{M} either $\models_{\mathfrak{M}} A$ or $\models_{\mathfrak{M}} \neg A$. So if \mathfrak{M} is a model of Γ then \mathfrak{M} is either a model of $\{A\} \cup \Gamma_1$ or of Γ_2 , by the definition of Γ_1, Γ_2 . \square

4.3.4. In our opinion the domain closure axiom is realistic in many applications (apart from the fact that every computer is a finite automaton). The lemmas 3.4-3.8 essentially depend on it. However, Theorem 4.3 remains valid when the domain closure axioms are removed from EQ . For, detected contradictions do not depend on the domain closure axioms and, conversely, consistency *with* the domain closure axiom trivially implies consistency *without* it.

4.4. Let us briefly discuss the semantical consequences of the conditions (2) and (3) from the previous theorem, since they may slightly deviate from the intended meaning of the knowledge base. It is possible that the consistency of Γ essentially depends on the valuation $f(c, \dots, c') = \text{undefined}$, i.e. that any valuation $f(c, \dots, c') = c_i$ ($1 \leq i \leq n$) would not yield a model for Γ . One could say that $f(c, \dots, c') = \text{undefined}$ possibly saves the expert system from inconsistencies by preventing production rules with occurrences of $f(c, \dots, c')$ in the antecedent from firing. Since such rules have obviously not been used in the inference, this may be considered an advantage. On the other hand, this may be considered a disadvantage in cases where $f(c, \dots, c') = \text{undefined}$ is not realistic (e.g. $\text{temperature}(\text{patient}) = \text{undefined}$). In these cases we suggest to add the appropriate unit clause $f(c, \dots, c') = c_i$ and to test consistency again.

4.5. Conceptually speaking it is not difficult to generalize the algorithm from 4.3 to universally quantified expert systems, although some care has to be taken in quantifying x in clauses containing literals of the form $\neg f(c, \dots, c') = x$. In these cases only restricted quantification of the form $\forall x \neq \text{undefined}$ is allowed. Of course, the algorithm from 4.3 can become very inefficient (from $P(c_0), P(c_1), \neg P(x_1) \vee \dots \vee \neg P(x_n) \vee Q(x_1, \dots, x_n)$, for instance, 2^n instances of Q are generated), so we suggest limited use of variables (or, preferably, the use of a more efficient algorithm).

4.6. **EXAMPLE.** We demonstrate the consistency test in action on the case of 1.2. Taking into account the object tree (sub-objects, single/multivalued attributes) the three rules yield the following clauses in shorthand:

$$\begin{aligned} & \neg C(p, cp) \vee P(p, c), \\ & \neg f_{abdp}(p) = \text{yes} \vee \neg f_{char}(p, pa) = co \vee \neg P(p, c), \\ & \neg C(p, ap) \vee f_{abdp}(p) = \text{yes}, \neg P(p, c) \vee f_{abdp}(p) = \text{yes}. \end{aligned}$$

When we add the unit clauses $C(p, cp)$ and $f_{char}(p, pa) = co$, the empty clause is easily derived and the algorithm decides to inconsistency. If we only add $C(p, cp)$, then the algorithm decides to consistency on the basis of the following (truth-)valuation:

$$\begin{aligned} & C(p, cp) = P(p, c) = \text{TRUE}, C(p, ap) = \text{FALSE}, \\ & f_{abdp}(p) = \text{yes}, f_{char}(p, pa) = \text{undefined}. \end{aligned}$$

5. EXPRESSIONS

5.1. The extension of the notion of indexed propositional expert system we discuss in this section aims at supporting arithmetic. We start with a real-life example, again taken from HEPAR. A similar example can be found in [BS], pg. 297.

greater_than <biochemistry, ((total_bili - direct_bili)/total_bili)*100,60>

This example shows an extension of the object-attribute-value representation by allowing *attribute expressions* instead of only attributes. A straightforward extension of the interpretation of $\langle o, a, v \rangle$ triples into many-sorted predicate logic yields the following atom in shorthand:

$$((tb(b) - db(b))/tb(b))*100 > 60$$

In order to capture attribute expressions, the notion of indexed propositional expert system has to be extended. To this end we add new function symbols, called *operators* and denoted by *op*, +, ..., to our language. Moreover the implicit axioms are extended with the positive diagram of each operator *op*, i.e. with all defining equations

$$op(c_i, \dots, c_j) = c_{f(i, \dots, j)},$$

where the constants are assumed to be of the sort that is required by the type of *op*. Here *f* does not belong to the language but denotes a function of indices describing the definition of *op*. Like = and <, operators are provided by the system. Let *OP* denote the set of axioms for operators. A typical example of an operator is addition of integers, denoted by + and written as an infix operator. From now on *IA* will denote $EQ \cup O \cup OP$.

5.2. We now extend the notion of *c*-atom as defined in Section 3 in the following way: a *c*-atom is a closed atom of the form $P(c, \dots, c')$ or of the form:

expression *relation symbol* *expression*

Here a relation symbol is one of the symbols <, =, > and an expression is a term built up from constants and terms of the form $f(c, \dots, c')$ by applying the operators (respecting types and sorts).

5.3. From now on indexed propositional expert systems as defined in 3.1 are understood to be extended in the sense of 5.1 and 5.2 above. In order to be able to carry over Theorem 4.3 we introduce the following definition.

DEFINITION. Assume every sort has a constant *undefined* as stated above in the conditions (2) and (3) of Theorem 4.3. An operator *op* is called *strict* if the implicit axioms for *op* satisfy the following condition:

$$op(c, \dots, c') = \text{undefined} \text{ if one or more of the arguments } c, \dots, c' \text{ equals } \text{undefined}.$$

Note that for a strict multiplication * of integers we have $0 * \text{undefined} = \text{undefined}$ and not $0 * \text{undefined} = 0$.

5.4. **THEOREM.** Let Γ be an axiomatization of an indexed propositional expert system satisfying the conditions (1)-(4) of Theorem 4.3 above. Assume moreover that the following conditions hold:

- (5) Every operator is strict and takes polynomial time.
- (6) No operators occur in the positive literals occurring in the explicit axioms of Γ .
- (7) All equations occurring in negative literals in the explicit axioms of Γ are of the form *expression* = *c*.

Then the consistency of Γ can be tested in polynomial time.

PROOF. Due to its general formulation, the algorithm from 4.3 carries immediately over to the present case. However, the argument demonstrating the correctness of this algorithm is slightly more complicated, and needs to be restated in some detail. We first note that the evaluation of expressions (such as the example in 5.1) may require more than one unit clause (both the values of $tb(b)$ and $db(b)$ are necessary). Furthermore, under the conditions (4) and (6) of the theorem, positive literals are of the form $P(c, \dots, c')$, $f(c, \dots, c') = c''$, $f(c, \dots, c') = g(d, \dots, d')$, or $c = c'$. The latter kind of literals evaluate immediately to *TRUE* or *FALSE* (by using *EQ*), so we may assume without loss of generality that they do not occur in the explicit axioms of Γ . Moreover, under condition (7) of the theorem, negative literals are essentially of the form $\neg P(c, \dots, c')$, $\neg \text{expression1} < \text{expression2}$ or $\neg \text{expression} = c$. By this more general form of literals involved, as well as the fact that $IA = EQ \cup O \cup OP$ instead of $IA = EQ \cup O$, the statements " $EQ \cup U_s$ is consistent" and " c -atom X is implied by $IA \cup U_s$ " have a more general meaning than in 4.3. For example, the former statement now also excludes $f(c) = 1, f(c) = g(d), g(d) = 2 \in U_s$, whereas the latter holds for X the c -atom from 5.1, with $-, /, *$ and $>$ having their usual meaning axiomatized in IA , and $tb(b) = 100, db(b) = 39 \in U_s$. The consistency in the THEN-part follows from the following (truth-) valuation:

$$P_i = \begin{cases} \text{TRUE} & \text{if } P_i \in U_s \\ \text{FALSE} & \text{otherwise} \end{cases}$$

$$t_j = \begin{cases} c & \text{if } EQ \cup U_s \models t_j = c \\ \text{undefined} & \text{otherwise} \end{cases}$$

Again we have that $EQ \cup U_s$ is consistent and that every clause not in U_s contains a negative literal L^- whose complement is not implied by $IA \cup U_s$. It can happen that $IA \cup U_s$ implies L^- (for example $L^- = \neg f(c) + f(c') < 0$, and $f(c) = f(c'), f(c') = 1 \in U_s$, with $+$ and $<$ their usual meaning). Then every clause containing L^- holds in any model of $IA \cup U_s$. If neither L^- nor L^+ is implied by $IA \cup U_s$, then we have either $L^+ = P(c, \dots, c') \notin U_s$, or L^+ contains a function term $f(c, \dots, c')$ such that $EQ \cup U_s$ does not imply $f(c, \dots, c') = c''$ for any constant c'' . By the properties of the constants *undefined* (conditions (2) and (3) of the theorem), the fact that every operator is strict (condition (5) of the theorem), as well as the syntactic restriction on negative literals (condition (7) of the theorem), the literals L^- mentioned above are *TRUE* in the (truth-) valuation. Hence every clause containing L^- is *TRUE* in the above (truth-) valuation, which constitutes (in the sense of 4.1) a model of the entire set of clauses. \square

REMARKS.

5.4.1. The intuitive idea behind the conditions (1), (4) and (6) is that the positive information, which can be seen as the driving force of the algorithm, should be definite. Omission of (1), (4) or (6) would allow the introduction of incomplete positive information of the form, respectively, $A^+ \vee B^+$, $f(c) < 2$, or $f(c) + f(c') = 3$. The other conditions (2), (3), (5) and (7) concern the simulation of partiality by means of the constants *undefined*, used for the circumvention of combinatorial explosions that would be caused by the search for consistent function values in a (possibly gigantic) product space of sorts. Thus all conditions are devoted to avoiding combinatorial explosions which would occur in the presence, implicit or explicit, of disjunctions of positive literals.

5.4.2. In Section 4 it is shown that the conditions (1)-(4) cannot be missed. It will be clear that condition (5) of Theorem 2.5, being quite natural, can also not be missed. We shall demonstrate furthermore that neither (6) nor (7) can be missed, by showing that the consistency test becomes NP-hard (see [GJ]) in cases in which (6), respectively (7), do not hold. This will be done in 5.4.3, respectively 5.4.4, by showing that certain disjunctions of positive literals are semantical consequences of indexed propositional expert systems with explicit axioms in Horn format. Then a similar argument as developed in 4.2, immediately yields the desired result.

5.4.3. As to (6), let β be a sort with exactly two constants, *TRUE* and *FALSE*, different from *undefined* ^{β} . Let *eq* be a strict equality operator of type $\beta \times \beta \rightarrow \beta$, i.e. *eq* yields *undefined* if one or both of its arguments equals *undefined* and otherwise *eq* yields *TRUE* if its arguments are equal, and *FALSE* else. It is easily seen that (with *f* of type $\dots \rightarrow \beta$)

$$EQ \cup OP \cup \{eq(f(c), f(c)) = TRUE\} \models f(c) = TRUE \vee f(c) = FALSE.$$

5.4.4. As to (7), let σ be a sort with exactly one constant 0 different from *undefined* ^{σ} . We trivially have (with *f* and *f'* of type $\dots \rightarrow \sigma$)

$$EQ \cup \{\neg f(c) = f'(c')\} \models f(c) = 0 \vee f'(c') = 0.$$

At first sight condition (7) seems to be a nasty restriction, since it excludes conditions of the form $f(c) = f'(c')$ in the rules. The reason is that =, being reflexive, does not propagate partiality properly (*undefined* = *undefined* evaluates to *TRUE*). There is, however, nothing against using strict equality operators in conditions, i.e. $eq^{\sigma \times \sigma \rightarrow \beta}(f(c), f'(c')) =_{\beta} TRUE$ instead of $f(c) =_{\sigma} f'(c')$.

CONCLUSION

We argued how to interpret a rule-based expert system as a many-sorted theory. Then the Tarski semantics yields a well-defined notion of consistency, and theorem proving techniques such as resolution can be used for testing consistency. The relevance for expert systems lies in the fact that inconsistencies make the conclusions of a knowledge-based system highly unreliable. For, what reason do we have to believe a conclusion if its negation could also be derived? We provided means for *detecting* inconsistencies, but did not discuss how to *deal* with them. Note that we considered an expert system as a fixed theory, whereas in practice the theory grows during the interaction with a user answering questions posed by the system (in this way the inconsistency of 1.2 was obtained). In our opinion only relatively consistent answers are to be accepted by the system. Starting with a consistent knowledge base we thus maintain consistency as an invariant of the interaction. Recent experiments with a PROLOG implementation of the consistency test and a medical knowledge base show that the models, produced by the consistency test at several moments during the interaction between the user and the system, are very informative.

REFERENCES

- [B] M. BEZEM, *Consistency of rule-based expert systems*. Proceedings of the 9-th Conference on Automated Deduction, Springer Lecture Notes in Computer Science 310, pp. 151-161, Springer-Verlag, Berlin (1987).
- [BS] B.G. BUCHANAN, E.H. SHORTLIFFE, *Rule-based expert systems: the Mycin experiments of the Stanford Heuristic Programming Project*. Addison-Wesley, Reading, Massachusetts (1984).
- [GJ] M.R. GAREY, D.S. JOHNSON, *Computers and intractability: a guide to the theory of NP-completeness*. Freeman, San Francisco, California (1979).
- [IL] T. IMIELINSKI, W. LIPSKI JR., *Incomplete information in relational databases*. Journal of the ACM 31, 4, pp. 761-791 (1984).
- [L] P.J.F. LUCAS, *Knowledge representation and inference in rule-based expert systems*. Report CS-R8613, Centre for Mathematics and Computer Science, Amsterdam (1986).
- [M] J.D. MONK, *Mathematical Logic*, Springer-Verlag, Berlin (1976).
- [R] J.A. ROBINSON, *Automatic deduction with hyper-resolution*. International Journal of Computer Mathematics 1, pp. 227-234 (1965).
- [Re] R. REITER, *Equality and domain closure in first order databases*. Journal of the ACM 27, 2, pp. 235-249 (1980).
- [W] C. WALTHER, *A many-sorted calculus based on resolution and paramodulation*. Pitman, London (1987).