Consiglio Nazionale delle Ricerche

Istituto di Elaborazione dell'Informazione

A terminological default logic

Fabrizio Sebastiani & Umberto Straccia

Nota Interna IEI-B4-54-1992

Novembre 1992

Via S. Maria, 46 - 56126 Pisa (Italy)
Phone : +39.(0)50.593400
Fax : +39.(0)50.554342
Telex : 590305 IEICNR I

A Terminological Default Logic

Fabrizio Sebastiani & Umberto Straccia
Istituto di Elaborazione dell'Informazione
Consiglio Nazionale delle Ricerche
Via S. Maria, 46 - 56126 Pisa (Italy)
E-mail : {fabrizio,straccia}@iei.pi.cnr.it
Phone : +39.50.593407
Fax : +39.50.554342

Abstract

Terminological Logics are knowledge representation formalisms of considerable applicative interest, as they are specifically oriented to the vast class of application domains that are describable by means of taxonomic organizations of complex objects. Although the field of terminological logics has lately been an active area of investigation, few researchers (if any) have addressed the problem of extending these logics with the ability to perform *default reasoning.* Such extensions would prove of paramount applicative value, as for many application domains a formalization by means of monotonic terminological logics may be accomplished only at the price of oversimplification. In this paper we show how we can effectively integrate terminological reasoning and default reasoning, yielding a *terminological default logic.* The kind of default reasoning we embed in our terminological logic is reminiscent of Reiter's Default Logic, but overcomes some of its drawbacks by subscribing to the "implicit" handling of exceptions typical of the Multiple Inheritance Networks with Exceptions proposed by Touretzky and others.

Keywords : Knowledge Representation, Knowledge Acquisition, Multiple Inheritance, Non-monotonic Reasoning

1 Introduction

Terminological Logics (TLs, variously known as *Frame Representation Languages* or *Concept Description Languages*) are knowledge representation formalisms of considerable applicative interest, as they are specifically oriented to the vast class of application domains that are describable by means of taxonomic organizations of complex objects.

Unlike better known logics (such as e.g. FOL), the primary syntactic expressions of TLs are *terms*, denoting monadic or dyadic relations on the domain of discourse. In general, the language of a TL consists of a number of term-forming operators by means of which one may build complex terms starting from a basic repertory of simple terms (*viz.* predicate symbols). By virtue of the semantics of such operators, a partial order is induced on the terms so defined, giving to a terminological KB the characteristic "hierarchical" (or taxonomic) structure of a directed acyclic graph.

The field of TLs has lately been an active area of research, with the attention of researchers especially focusing on the investigation of their logical and computational properties. Nevertheless, few researchers (if any) have addressed the problem of extending these logics with the ability to perform *default reasoning*, a kind of non-monotonic reasoning that is to be applied whenever the rules involved allow for *exceptions.*

Non-monotonic reasoning has been formally addressed in various ways, leading to the development of a variety of formalisms, most of which belong to the offspring of Doyle and McDermott's Nonmonotonic Logic [3], Reiter's Default Logic [13] and McCarthy's Circumscription [8]. Each of these formalisms may be seen as extending FOL with non-monotonic reasoning capabilities of some kind. Given that TLs may be viewed as (pragmatically and computationally interesting subsets of FOL, one might be led to think that a simple integration of default and terminological reasoning could be obtained by simply considering one of the non-monotonic formalisms mentioned above and restricting it to deal with the chosen TL, rather than with FOL *tout court.* Unfortunately, these non-monotonic formalisms, besides having unattractive metalogical and computational properties, suffer from a problem that hinders their use in KR contexts requiring that KB construction be accomplished in an incremental fashion. We call this problem the *Exceptions Explicitation Problem* (EEP).

Incrementality of KB construction is an asset of KB management systems that hardly needs to be argued for. Large KBs are the result of an evolutionary process, both because knowledge entry is a time-consuming process and because knowledge may simply become available at later stages of the process. Also, when a large KB is built by this "stepwise refinement" process, it is highly desirable that the refinement consists in the plain, piecemeal *addition* of new knowledge chunks rather than in a time-consuming *revision* (with possibly ensuing deletion) of pre-existing chunks.

In the full paper [15] we discuss, by means of concrete examples, the EEP and how it manifests itself, for example, in the context of Nonmonotonic Logic (NML) (to this respect, other formalisms such as Default Logic and Circumscription behave in a completely analogous way); we also discuss how the addition of an NML formula to a KB calls for a revision of the pre-existing KB that may in general require repeated calls to the NML theorem prover, an endeavour that we deem absurd, given the intractability and undecidability of NML. The net effect is that, unless the construction of the KB is realized in a completely static (non incremental) way, the problem of KB construction in NML is in fact an *unsolvable* problem.

While the general non-monotonic formalisms mentioned above are affected by the EEP, this is not true of the formalisms for *Multiple Inheritance Networks with Exceptions* (MINEs), a popular, albeit less general, class of non-monotonic KR languages oriented to the representation of taxonomic knowledge (see e.g. [16]). Such languages are less expressive than the more general non-monotonic formalisms mentioned above, in the sense that they only allow for monadic predicate symbols, a limited use of negation and no disjunction at all. For our purposes, it is also essential to observe that their monotonic fragment is far less expressive than TLs as, having no term constructors in their syntactic apparatus, they only allow the definition of taxonomies of simple predicate symbols.

MINEs do not suffer from the EEP because they implement an *implicit* handling of exceptions by exploiting the partial ordering given by the taxonomy: as a first approximation we can say that, in case of conflicts, a "default rule" $a \rightarrow b$ is preferred to another default rule $c \not\rightarrow b$ if the precondition of the first precedes the precondition of the second in the ordering. In other words, the implicit handling of exceptions obeys the so-called *specialization principle*, according to which "conflicts" are to be solved by preferring the properties belonging to a subclass over those belonging to a superclass.

In this paper we will show how we can effectively extend TLs in such a way that they allow a brand of default reasoning that obeys the specialization principle, thus creating a logic that combines the tools for describing taxonomic organizations of complex objects which are typical of TLs, the ability to describe default information which is typical of general nonmonotonic formalisms, and the incrementality in KB construction which is typical of MINEs. Such an endeavour constitutes perhaps the first completely formal realization of the notion of "frame" as originally proposed by Minsky [9], a notion that was intended to describe a highly structured aggregate of knowledge allowing the description of "prototypical" knowledge and resulting in KBs of "taxonomic" form. We will call our logic $\mathcal{TDL}^-$ (Terminological Default Logic—the "-" superscript distinguishes it from an earlier version).

This paper is organized as follows. In Section 2 we formally introduce the syntax and the semantics of the monotonic fragment of $\mathcal{TDL}^-$. In Section 3 we deal with the non-monotonic part of $\mathcal{TDL}^-$, describing the notion of "extension" of a set of $\mathcal{TDL}^-$ formulae (i.e. the set of conclusions that may be derived from these formulae) and some of its properties. In Section 4 we discuss an algorithm that computes an extension of a set of $\mathcal{TDL}^-$ formulae (when it exists), and discuss the issue of the computational complexity of $\mathcal{TDL}^-$. For reasons of space, the proofs of the propositions and theorems stated in this paper are omitted; see [15] also for a more detailed account of the logical and computational properties of $\mathcal{TDL}^-$. Section 5 concludes.

2 The monotonic fragment of $\mathcal{TDL}^-$

The $\mathcal{TDL}^-$ logic, like many other TLs, allows the specification of three fundamental types of terms: *frames*, *slots* and *individual constants*. Frames (also known as *concepts*) are terms denoting sets of individuals, and are, so to speak, the first-class citizens of $\mathcal{TDL}^-$. Slots (also known as *roles*) are terms denoting binary relations between individuals; their function is to allow the specification of structural constituents of frames. Individual constants denote individuals of the domain of discourse. For example, the basic repertory of simple terms (called *atoms*) that are used in order to build more complex terms might contain the frame `Polygon`, denoting the set of polygons, and the slot `Side`, denoting all those pairs of individuals $\langle x, y\rangle$ such that y is one of the sides of x; this would allow the definition of more complex frames, such as e.g. the term $\forall$`Side.Polygon`, denoting the set of those individuals whose sides are all polygons (i.e. the set of polyhedra), and to subsequently define other frames by using those defined before.

In order to introduce the syntax of $\mathcal{TDL}^-$ we will need three disjoint alphabets: an alphabet I of individual constants (with metavariables i, i_1, i_2, ...), an alphabet MP of monadic predicate symbols (with metavariables M, M_1, M_2, ...) and an alphabet DP of dyadic predicate symbols (with metavariables D, D_1, D_2, ...). The syntax of $\mathcal{TDL}^-$ is specified by the following definition.

Definition 1 *A* frame *in $\mathcal{TDL}^-$ is defined by the following syntax:*

$$\begin{array}{rcl} F_1, F_2 & \rightarrow & F_1 \sqcap F_2 \mid M \mid \neg M \mid \forall S.F_1 \mid \bot \mid \top \\ S & \rightarrow & D \end{array}$$

We will use metavariables $F, F_1, F_2, \ldots$ ranging on frames and metavariables $S, S_1, S_2, \ldots$ ranging on slots.

Let us now switch to the formal semantics of $\mathcal{TDL}^-$ frames. The meaning of the linguistic constructs introduced above may be given in terms of the notion of extension function.

Definition 2 *An* interpretation $\mathcal{I}$ *over a nonempty set of individuals $\mathcal{D}$ is a function that maps elements of MP into subsets of $\mathcal{D}$, elements of DP into subsets of $\mathcal{D}\times\mathcal{D}$, and elements of I into elements of $\mathcal{D}$ such that $\mathcal{I}(i_1) \neq \mathcal{I}(i_2)$ whenever $i_1 \neq i_2$. We will say that $\mathcal{I}$ is an* extension function *iff $\mathcal{I}(\bot) = \emptyset$, $\mathcal{I}(\top) = \mathcal{D}$, $\mathcal{I}(\neg M) = \mathcal{D} \setminus \mathcal{I}(M)$, $\mathcal{I}(F_1 \sqcap F_2) = \mathcal{I}(F_1) \cap \mathcal{I}(F_2)$, $\mathcal{I}(\forall S.F) = \{x \in \mathcal{D} \mid \forall y : \langle x, y\rangle \in \mathcal{I}(S) \Rightarrow y \in \mathcal{I}(F)\}$.* ∎

We next introduce a feature of the language that allows us to associate names to complex frames, with the net effect that we will be able to define new frames using names instead of the corresponding complex frames.

Definition 3 *A* naming *is an expression of the form $M \doteq F$ or of the form $M \dot{<} F$, where F is a frame and M an element of MP. An extension function $\mathcal{I}$ over a nonempty domain $\mathcal{D}$* satisfies *a naming δ iff $\mathcal{I}(M) = \mathcal{I}(F)$ if $\delta = M \doteq F$, and $\mathcal{I}(M) \subseteq \mathcal{I}(F)$ if $\delta = M \dot{<} F$.* ∎

Namings of type $M \doteq F$ define necessary and sufficient conditions for an individual to be in the extension of M, while the conditions defined by namings of type $M \dot{<} F$ are necessary but not sufficient.

Our language also allows the expression of assertions, stating that individual constants are instances of frames and pairs of individual constants are instances of slots.

Definition 4 *An* assertion *is an expression of the form $i_1{:}F$ or $\langle i_1, i_2\rangle{:}S$, where i_1, i_2 are individual constants, F is a frame and S is a slot. An extension function $\mathcal{I}$ over a nonempty domain $\mathcal{D}$* satisfies an assertion *α iff $\mathcal{I}(i_1) \in \mathcal{I}(F)$ if $\alpha = i_1{:}F$ and $\langle \mathcal{I}(i_1), \mathcal{I}(i_2)\rangle \in \mathcal{I}(S)$ if $\alpha = \langle i_1, i_2\rangle{:}S$.* ∎

The notion of *satisfiability* of a *T-set* Ω (i.e. of a set of namings and assertions), and that of a *model* of a T-set Ω may be defined in the obvious way (see [15]).

Definition 5 *Let Ω be a satisfiable T-set, and let M_1, M_2 be two elements of MP. We say that M_1 is* subsumed *by M_2 in Ω (written $M_1 \preceq_\Omega M_2$) iff for every model $\mathcal{I}$ of Ω it is true that $\mathcal{I}(M_1) \subseteq \mathcal{I}(M_2)$. We say that Ω* logically implies *an assertion α (written $\Omega \models \alpha$) iff α is satisfied by all models of Ω.* ∎

The following definition formalizes what namings and assertions follow from a T-set Ω.

Definition 6 *Let Ω be a satisfiable T-set. The* transitive closure *of Ω (written $TC(\Omega)$) is the set $\Omega \cup \{\alpha \mid \Omega \models \alpha\}$.* ∎

Observe that $\models$ is defined for assertions only: this means that a set Ω differs from $TC(\Omega)$ only in the assertions it contains. It is easy to show that TC is monotonic (i.e. if $\Omega \subseteq \Omega'$, then $TC(\Omega) \subseteq TC(\Omega')$), and that TC is in fact a closure (i.e. $TC(\Omega) = TC(TC(\Omega))$).

3 Default reasoning in $\mathcal{TDL}^-$

Up to now we have described the monotonic fragment of $\mathcal{TDL}^-$. Let us now discuss the addition of non-monotonic features.

Definition 7 *A* default *is an expression of the form $M \mapsto S.F$, where M is an element of MP, S is a slot and F is a frame.* ∎

Informally, $M \mapsto S.F$ means: "if i is an M such that i' is an S-filler of i and the assumption that i' is a F does not lead to a contradiction, assume it". For example, the default $\texttt{IU} \mapsto \texttt{FM.I}$ means: "if i is an Italian university and i' is a faculty member of it and the assumption that i' is Italian does not lead to a contradiction (i.e. we do not know that he is not an Italian), assume it".

The particular syntax we have chosen for defaults is due to the following reasons:

1. an analysis of the literature concerning the interaction between frames and default knowledge, ranging from the more informal and impressionistic proposals (such as those of e.g. [9,14]) to more formally justified ones [2,11], reveals that default rules with consequents in a "slot-filler" form have always been identified as the most natural way in which to convey default frame-like knowledge;

2. this type of default rules are sufficient to highlight the problems resulting from the interaction between default knowledge and terminological knowledge; on the other hand, the extension to other forms of defaults (such as e.g. those involving numeric restrictions) is conceptually easy, but does not teach us much with respect to the issue of the integration between default and terminological reasoning.

We may now define what we mean by a $\mathcal{TDL}^-$ theory.

Definition 8 *A* $\mathcal{TDL}^-$ theory *is a pair* $\mathcal{T} = \langle \Psi, \Delta \rangle$*, where* Ψ *is a satisfiable finite T-set, and* Δ *is a finite set of defaults.* ∎

We are now able to define what the *extensions* of a $\mathcal{TDL}^-$ theory are. Informally, by the term "extension" we mean the set of assertions and namings that we can reasonably believe to be true as a consequence of the theory. For instance, if we knew that the University of Bellosguardo is an Italian university, that Professor Dolcevita is a faculty member thereof, and that the faculty members of Italian universities are typically Italian, we would like to conclude (formally: to be included in the corresponding extension) that Dolcevita is an Italian.

Our definition of "extension" is similar to the one given by Reiter for Default Logic, i.e. an extension is a fixpoint of a consequence relation. However, unlike in Default Logic, the specialization principle is "wired" in our definition: in the presence of conflicting defaults, the one with the more specific premise will be preferred. For instance, suppose that, besides the fact that the faculty members of Italian universities are typically Italian, we also knew that the faculty members of South Tyrolean universities are typically not Italian; knowing that South Tyrolean universities are Italian universities[1], that the University of Pflunders is a South Tyrolean university, and that Professor Katzenjammer is a faculty member thereof, we will be able to derive, as desired, that Katzenjammer is not Italian. Such a conclusion could not be drawn if we simply confined ourselves to employing the terminological subset of Default Logic (or, for that matter, of any general non-monotonic formalism): the specializazion principle embodied in Definition 9 plays a critical role in the inferential behaviour displayed by our formalism.

Definition 9 *Let $\mathcal{T} = \langle \Psi, \Delta \rangle$ be a $\mathcal{TDL}^-$ theory. Let Γ be an operator such that, for any satisfiable T-set Ω, $\Gamma(\Omega, \mathcal{T})$ is the smallest satisfiable T-set satisfying the following closure conditions:*

1. $\Psi \subseteq \Gamma(\Omega, \mathcal{T})$;
2. $\Gamma(\Omega, \mathcal{T}) = TC(\Gamma(\Omega, \mathcal{T}))$;
3. *for all defaults $M_1 \mapsto S.F_1 \in \Delta$, for all assertions $i_1{:}M_1 \in \Gamma(\Omega, \mathcal{T})$ such that $\langle i_1, i_2 \rangle{:}S \in \Gamma(\Omega, \mathcal{T})$, it happens that $i_2{:}F_1 \in \Gamma(\Omega, \mathcal{T})$, unless there exists an atom M_2 such that*
   (a) $i_1{:}M_2 \in \Omega$ *and* $M_2 \preceq_\Omega M_1$;
   (b) $M_2 \mapsto S.F_2 \in \Delta$;
   (c) $\Omega \cup \{i_2{:}F_1 \sqcap F_2\}$ *is unsatisfiable.*

A satisfiable T-set $\mathcal{E}$ is an extension *of the $\mathcal{TDL}^-$ theory $\mathcal{T}$ iff $\mathcal{E} = \Gamma(\mathcal{E}, \mathcal{T})$, i.e. iff $\mathcal{E}$ is a fixpoint of the operator Γ.* ∎

Conditions 1 and 2 are obviously to be satisfied if we want "extensions" to be "sets of conclusions" according to the sense generally accepted in KR. Condition 3 embodies the specialization principle: if a default $M_1 \mapsto S.F_1$ is "applicable" and (Conditions 3a,3b,3c) there is no evidence contradicting the conclusion of the default (i.e. i_1 does not belong to any subclass of M_1 which

[1]South Tyrol is, in fact, a German-speaking region of Italy.

is the premise of a default whose conclusion would be inconsistent with F_1), then the default may be safely applied and the conclusion drawn.

We now consider an example to show how Definition 9 works, and, in particular, how $\mathcal{TDL}^-$ employs an implicit handling of exceptions.

Example 1 Let $\mathcal{T} = \langle \Psi, \Delta \rangle$ be the $\mathcal{TDL}^-$ theory that formalizes our "Professors" example, with $\Psi = \{$`b:IU`, $\langle$`b,d`$\rangle$`:FM`, `p:STU`, $\langle$`p,k`$\rangle$`:FM`, `STU` $\lessdot$ `IU`$\}$ and $\Delta = \{\mathtt{IU}{\mapsto}\mathtt{FM.I}, \mathtt{STU}{\mapsto}\mathtt{FM.\neg I}\}$. Let $\mathcal{E} = TC(\Psi \cup \{\mathtt{k{:}\neg I}, \mathtt{d{:}I}\})$ and $\Gamma(\mathcal{E}, \mathcal{T}) = \mathcal{E}$. It is not hard to show (see [15]) that $\Gamma(\mathcal{E}, \mathcal{T})$ satisfies the conditions of Definition 9; therefore $\mathcal{E}$ is an extension of $\mathcal{T}$.

It is important to observe that the same example may be formalized, for example, in Nonmonotonic Logic, only at the price of a cumbersome operation of "exceptions explicitation", i.e. by imposing the following set of axioms.

$$\forall x \forall y \; IU(x) \wedge FM(x,y) \wedge M[I(y) \wedge \neg STU(x)] \Rightarrow I(y) \tag{1}$$
$$\forall x \; STU(x) \Rightarrow IU(x) \tag{2}$$
$$\forall x \forall y \; STU(x) \wedge FM(x,y) \wedge M[\neg I(y)] \Rightarrow \neg I(y) \tag{3}$$
$$IU(b) \wedge FM(b,d) \wedge STU(p) \wedge FM(p,k) \tag{4}$$

As Axiom 1 shows, in NML we must make explicit the fact that a South Tyrolean university is an exceptional Italian university relatively to the citizenship of its faculty members; in $\mathcal{TDL}^-$ this is not necessary, and, as hinted in Section 1, this allows KB update to be completely *additive*. ∎

In [15] we use Example 1 to show that $\mathcal{TDL}^-$ is in fact non-monotonic.

We go on to discuss some properties of the notion of extension as formalized in Definition 9. The following proposition parallels the one given by Reiter for Default Logic, stating that extensions are "maximal" sets.

Proposition 1 *Let $\mathcal{T} = \langle \Psi, \Delta \rangle$ be a $\mathcal{TDL}^-$ theory, and let $\mathcal{E}_1$ and $\mathcal{E}_2$ be extensions of $\mathcal{T}$. If $\mathcal{E}_1 \subseteq \mathcal{E}_2$, then $\mathcal{E}_1 = \mathcal{E}_2$.* ∎

Similarly to what happens in most non-monotonic formalisms, some $\mathcal{TDL}^-$ theories have more than one extension; in particular, the number of extensions can be exponential with respect to the size of the $\mathcal{TDL}^-$ theory, as the following proposition shows.

Proposition 2 *There exists a $\mathcal{TDL}^-$ theory $\mathcal{T}$ such that the number of extensions of $\mathcal{T}$ is exponential with respect to the size of $\mathcal{T}$.* ∎

This proposition is proven by showing that there exists a $\mathcal{TDL}^-$ theory $\mathcal{T}$ that contains $O(|\mathcal{T}|^2)$ "Nixon Diamonds", which leads to the existence of $O(2^{|\mathcal{T}|^2})$ extensions. Although the number of extensions can be very large in the worst case, this needs not be the case in actual KBs. Furthermore, this exponential number of extensions is not a characteristic of $\mathcal{TDL}^-$ itself, but is common to all standard non-monotonic formalisms: the "Multiple Nixon Diamond" can be easily formulated in these formalisms, giving rise to the same phenomenon.

Unfortunately, some $\mathcal{TDL}^-$ theories may not have extensions.

Proposition 3 *There exists a $\mathcal{TDL}^-$ theory with no extensions.* ∎

This is proven by showing that the sample theory $\mathcal{T} = \langle \Psi, \Delta \rangle$, with $\Psi = \{$a:A, $\langle$a,a$\rangle$:S,A$\doteq$B$\sqcap$C, D$\doteq$B$\sqcap$C$\sqcap$E$\}$ and $\Delta = \{$A$\mapsto$S.E, D$\mapsto$S.$\perp\}$, has no extensions. In [15] we argue that it would not be easy to find sublanguages of $\mathcal{TDL}^-$ such that the existence of at least one extension is always guaranteed: in fact, removing from $\mathcal{TDL}^-$ the causes that are responsible for the non-existence of an extension for theory $\mathcal{T}$ would dramatically curtail the expressive power of the language itself.

4 Computing an extension of a $\mathcal{TDL}^-$ theory

In this section we will discuss the properties of the EXT (nondeterministic) algorithm that computes (when it exists) an extension of a $\mathcal{TDL}^-$ theory.

EXT is heavily dependent on the decision of the monotonic fragment of $\mathcal{TDL}^-$, i.e. of the $\preceq_\Omega$ and $\models$ relations. It is well-known (see [6]) that in most TLs (and the monotonic fragment of $\mathcal{TDL}^-$ is no exception), deciding $\preceq_\Omega$ can be reduced to the decision of $\models$, and that the decision of $\models$ can in turn be reduced to deciding unsatisfiability[2].

There exists a well-known technique, based on constraint propagation (see [6]), for deciding unsatisfiability in TLs. By using this technique, it may be shown that it is decidable whether a finite and "acyclic" T-set (i.e. a T-set that contains no namings in which the *definiendum* is defined in terms of

[2]The unsatisfiability problem is the problem of deciding if a T-set Ω is unsatisfiable.

itself) is unsatisfiable. By considering acyclic $\mathcal{TDL}^-$ theories only, we can profitably exploit this result.

Definition 10 *Let $\mathcal{T} = \langle \Psi, \Delta \rangle$ be a $\mathcal{TDL}^-$ theory. Then $\mathcal{T}$ is* acyclic *iff $\Phi = \{\tau | \tau$ is a naming in $\Psi\} \cup \{M \lessdot F{:}M \mapsto S.F \in \Delta\}$ is acyclic.* ∎

In [15] we present the EXT algorithm for computing an extension of an acyclic $\mathcal{TDL}^-$ theory by a series of successive approximations Ω_i. If two successive approximations are the same set Ω_n, the algorithm is said to *converge*, and Ω_n is a finite and acyclic T-set such that $TC(\Omega_n)$ is an extension of the theory. EXT contains a loop inside which a nondeterministic choice is made of which default to consider for expansion. Generality requires this nondeterminism, since $\mathcal{TDL}^-$ theories need not have unique extensions.

The following correctness and completeness theorem states that all and only the extensions of a $\mathcal{TDL}^-$ theory $\mathcal{T}$ can be computed by the algorithm.

Theorem 1 *Let $\mathcal{T} = \langle \Psi, \Delta \rangle$ be an acyclic $\mathcal{TDL}^-$ theory. $\mathcal{E}$ is an extension of $\mathcal{T}$ iff the application of the* EXT *algorithm to $\mathcal{T}$ has a converging computation such that $\Omega_n = \Omega_{n-1}$ and $TC(\Omega_n) = \mathcal{E}$.* ∎

The following is an example of a non-converging computation.

Example 2 Consider the acyclic $\mathcal{TDL}^-$ theory $\mathcal{T}$ of Proposition 3. It turns out that each approximation Ω_i is such that $\Omega_{2k} = \Psi$ and $\Omega_{2k+1} = \Psi \cup \{\mathrm{a{:}E}\}$, for each $k \geq 0$. Therefore $\Omega_{2k} \neq \Omega_{2k+1}$ for each $k \geq 0$, and the computation never stops. ∎

The next corollary follows from Theorem 1.

Corollary 1 *The set of extensions of an acyclic $\mathcal{TDL}^-$ theory is recursively enumerable.* ∎

Finally, we discuss some issues related to the computational complexity of $\mathcal{TDL}^-$. To this respect, it is fundamental to observe that, if we add an assertion α to Ω, $\preceq_\Omega$ does not change, as the following proposition states.

Proposition 4 *Let Ω be a T-set in $\mathcal{TDL}^-$, M_1 and M_2 two elements of MP, and α an assertion. Then $M_1 \preceq_{\Omega \cup \{\alpha\}} M_2$ iff $M_1 \preceq_\Omega M_2$.* ∎

This also means that, if we want to compute an extension of a $\mathcal{TDL}^-$ theory $\mathcal{T} = \langle \Psi, \Delta \rangle$, it suffices to compute $\preceq_\Psi$ at the beginning of the computation, once for all. Unfortunately, deciding $\preceq_\Psi$ is co-NP-complete, as shown in [12]. Luckily, this is a worst case behaviour that seldom occurs: if we make some reasonable assumptions (see [12]) on the form of Ψ, it can be shown [15] that checking if $M_1 \preceq_\Psi M_2$ holds is computable in $O(s \log s)$, where $s = |\Psi|$.

5 Conclusion

In this paper we have shown how we can extend terminological logics in such a way that they allow a brand of default reasoning that obeys the specialization principle, thus creating a formalism (which we have dubbed $\mathcal{TDL}^-$) that combines the tools for describing taxonomic organizations of complex objects which are typical of TLs, the ability to describe default information which is typical of general nonmonotonic formalisms, and the incrementality in KB construction which is typical of MINEs.

This has been obtained by relying on the notion of "extension of a $\mathcal{TDL}^-$ theory", a notion that has been defined in the style pioneered by Reiter in his Default Logic, i.e. as a fixpoint of a consequence relation. We have also studied a number of properties of $\mathcal{TDL}^-$ related to issues such as the existence and the uniqueness of extensions, and the complexity of $\mathcal{TDL}^-$ reasoning.

The language of $\mathcal{TDL}^-$ has been designed with the aim of providing a minimal framework allowing to study the interaction of terminological and default information in a meaningful way. Quite obviously, extensions to this framework may be conceived that enable the expression of default information of a nature different from the one considered here.

References

[1] Ronald J. Brachman and Hector J. Levesque, editors. *Readings in knowledge representation.* Morgan Kaufmann, Los Altos, CA, 1985.

[2] Gerhard Brewka. The logic of inheritance in frame systems. In *Proceedings of IJCAI-87, 10th International Joint Conference on Artificial Intelligence*, pages 483–488, Milano, Italy, 1987.

[3] Jon Doyle and Drew McDermott. Nonmonotonic logic I. *Artificial Intelligence*, 13:41–72, 1980. [a] Appears also in [4], pp. 111–126.

[4] Matthew L. Ginsberg, editor. *Readings in nonmonotonic reasoning.* Morgan Kaufmann, Los Altos, CA, 1987.

[5] John Haugeland, editor. *Mind design.* The MIT Press, Cambridge, MA, 1981.

[6] Bernhard Hollunder. Hybrid inferences in KL-ONE-based knowledge representation systems. In *Proceedings of the 14th German Workshop on Artificial Intelligence*, pages 38–47, Eringerfeld, FRG, 1990.

[7] Vladimir Lifschitz, editor. *Programs with common sense - Papers by John McCarthy.* Ablex, Norwood, NJ, 1990.

[8] John McCarthy. Circumscription - a form of nonmonotonic reasoning. *Artificial Intelligence*, 13:27–39, 1980. [a] Appears also in [4], pp. 145–151. [b] Appears also in [7], pp. 142–155.

[9] Marvin Minsky. A framework for representing knowledge. In Patrick J. Winston, editor, *The psychology of computer vision*, pages 211–277, McGraw-Hill, New York, NY, 1975. [a] An extended version appears also in [1], pp. 245–262, and in [5], pp. 95–128.

[10] Robert C. Moore. Semantical considerations on nonmonotonic logic. *Artificial Intelligence*, 25:75–94, 1985. [a] Appears also in [4], pp. 127–136.

[11] Robert A. Nado and Richard E. Fikes. Semantically sound inheritance for a formally defined frame language with defaults. In *Proceedings of AAAI-87, 6th Conference of the American Association for Artificial Intelligence*, pages 443–448, Seattle, WA, 1987.

[12] Bernhard Nebel. Terminological reasoning is inherently intractable. *Artificial Intelligence*, 43:235–249, 1990.

[13] Raymond Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980. [a] Appears also in [4], pp. 68–93.

[14] R. Bruce Roberts and Ira P. Goldstein. *The FRL manual.* Technical Report 409, The Artificial Intelligence Laboratory, MIT, Cambridge, MA, 1977.

[15] Fabrizio Sebastiani and Umberto Straccia. *A terminological default logic (extended report).* Technical Report, Istituto di Elaborazione dell'Informazione, Consiglio Nazionale delle Ricerche, Pisa, Italy, 1992. Forthcoming.

[16] David S. Touretzky, John F. Horty, and Richmond H. Thomason. A clash of intuitions: the current state of nonmonotonic multiple inheritance systems. In *Proceedings of IJCAI-87, 10th International Joint Conference on Artificial Intelligence*, pages 476–482, Milano, Italy, 1987.

Towards a Relevance Logic of Information Retrieval

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

Abstract

We propose relevance logics as a possible foundation for information retrieval and indicate research directions to this end.

1 Information retrieval and relevance

As recently pointed out by van Rijsbergen [13], the discipline of Information Retrieval (IR) is *scientific*, in the sense that its underlying methodology is heavily dependent on the experimental method: the design and implementation of an IR system is followed by a phase of testing experiments in which

1. the system is confronted with a series of queries formulated by a person taking part in the experiment and playing the part of the user;
2. the system evaluates these queries against a body of information, thus retrieving the information that, according to the retrieval strategy underlying the system, is deemed relevant to the "user" queries;
3. the user judges the relevance of the information the system has retrieved (in terms of *precision* and *recall*), presumably bringing to bear his/her own understanding of what the queries were intended to mean.

This cycle is performed repeatedly, with a number of sample queries being proposed to the system, and a corresponding number of "relevance judgements" being made by the user; these combined "relevance judgements" are the only parameters that, in the end, concur in finally assessing the adequacy of the system.

Relevance is then the most critical parameter in the evaluation of IR systems; bulding systems that perform the retrieval of relevant information only may thus be considered the most important issue in IR. Unfortunately, the primacy of relevance in the whole IR discipline is also the primary cause that has hindered, up to now, the development of a *theory* of IR. In fact, relevance is not a formally and clearly defined notion; what relevance is, in other words, is defined by the user from time to time and from experiment to experiment, and is then heavily dependent on judgements where highly subjective and scarcely reproducible factors are brought to bear. The very possibility of a theory of IR is then dependent on the possibility of giving a *formal* definition of what relevance is, a definition capable of abstracting from the subjective and contingent factors inherent in the *operational* view of relevance described above.

2 Mathematical logic, conditionals and relevance

Some recent works (see e.g. [13]) have thus addressed the foundational problem of IR by trying to give a formal notion of relevance based on mathematical logic. These researches have shown how the relevance of a document d to a query q may naturally be understood in terms of a *conditional* (sometimes also called an *implication*) $d \rightarrow q$, where the "$\rightarrow$" symbol is the particular conditional notion formalized by a given logic. The foundational problem has then become the problem of singling out the logic (or those logics), among the ones that deal with conditional notions at all, whose conditional takes into account "relevance" as a critical factor.

The history of logic has seen a flurry of logics motivated by the need to give a natural account of the conditional. Classical logic itself possesses a well-known conditional notion, *material implication* (denoted by the symbol "$\supset$"). However, material implication has often been criticized, on the account that it licenses paradoxical sentences as theorems of the pure calculus; for instance, the sentence $(a \supset (b \supset a))$ (asserting that a true proposition is implied by any proposition) is a theorem of classical logic, a state of affairs that is questionable at best. The driving force behind the development of modal logic was exactly the desire to provide a conditional notion that did not fall short of the paradoxes of material implication. Unfortunately, the conditional notion formalized by modal logic (called *strict implication*, and denoted by the symbol "$\leadsto$"), although solving the paradoxes of material implication, suffers from paradoxes of its own, making it a no more viable alternative than classical logic; for instance, the sentence $(La \supset (b \leadsto a))$ (asserting that a necessary proposition is strictly implied by any proposition) is a theorem of modal logic, a no less paradoxical fact than the one discussed above.

From our standpoint of would-be IR theorists, it is interesting to note that some of the paradoxical sentences belonging to classical and modal logics are actually conditional sentences that suffer from *fallacies of relevance*: in other words, they are theorems of the given logic *even if their premise is not relevant to their conclusion.* For instance, the fact that $(a \supset (b \supset a))$ (resp. $(La \supset (b \leadsto a)))$ is valid in classical (resp. modal) logic should strike one as peculiar, in that in any of these cases the fact that b holds does not have any "relevance" to the fact that a holds! For example, although a man asserting that "if Fermat's last conjecture is true, then Rome is the capital of Italy" would assert a true sentence according to most logicians, he would simply utter nonsense according to the man in the street! Indeed, the tradition of modern logic that has departed from Frege and Russel-Whitehead, seems to have proceeded without paying too much attention to relevance, somehow implicitly assuming that this notion belonged more to the realm of rhetorics than the one of logic.

3 Relevance logics

Among the first to take such a stand, Nelson [10] has argued that, in order for any conditional notion "$\rightarrow$" to be adequate, a sentence such as $a \rightarrow b$ should be valid only if there be "some connection of meaning between a and b". To the surprise of many orthodox logicians, the idea of a "connection of meaning between a and b" (or, more generally, the idea of a being *relevant* to b) has been shown to be amenable to formal treatment by a number of logicians who have defined a class of logical calculi called *relevance* or *relevant logics* [1,5].

Relevance logics attempt to formalize a conditional notion in which relevance is a primary concern. By doing this, they challenge classical logic in a number of ways [6], i.e. by introducing a new, non truth-functional connective (denoted by "$\rightarrow$") into the syntactic apparatus of classical logic, by rejecting some classical rules of inference for classical connectives, and by changing the notion of validity itself by "wiring" into it considerations of relevance.

As with modal logics, there are many relevance logics; some of them are ordered with respect to expressive power, while some of them are incommensurable with respect to this dimension; more importantly, different relevance logics formalize a different notion of relevance.

4 Towards a relevance logic of information retrieval

The aim of this document is to suggest that *relevance logic is a promising candidate to become the logic of information retrieval.* In fact, even a brief analysis of the motivations put forth by relevance logicians and by IR theorists, respectively, indicates a surprising coincidence of underlying tenets and purposes. Therefore, it seems just natural to think that, if we view retrieval as essentially consisting of a disguised form of logical inference [13], IR and relevance logic might constitute the engineering side and the theoretical side of the same coin. As a consequence, it seems just natural to investigate the relationships between the two disciplines, with the aim of establishing relevance logic as a foundation of IR.

This investigation should be structured according to the points illustrated in the following subsections.

4.1 Pre-theoretic foundations

Before diving into formal complexities, an *empirical* study should be undertaken, with the principal aim of assessing whether there is an actual and deep coincidence of views between logicians and IR theorists that attempt to establish the notion of relevance as the cornerstone of their theories. This investigation may largely proceed without concerns about mathematical details. Among other things, it should compare the different notions of relevance formalized by the different relevance logics, and look for the ones whose underlying motivations comply most with the principles of information retrieval.

We think it would be promising to start the investigation from a specific relevance logic that seems to comply with some of the requirements of the IR world: the logic $\mathbf{E_{fde}}$ of *tautological entailments* [4]. $\mathbf{E_{fde}}$ is the fragment of the relevance logics **E** and **R** (called the logic of *entailment* and the logic of *relevant implication*, respectively) that deals with *first degree entailments* only, i.e. pairs of propositional (classical) formulae separated by one "$\rightarrow$" symbol. This logic seems well suited to formalize a state of affairs in which both document and query have a boolean representation, and in which the relevance of one to the other is the parameter of interest.

4.2 Formal foundations

A *formal* study should then be undertaken, aimed at investigating the role of the logic which has supposedly been chosen in the preceding phase, and extending it in various ways. This study should address the following aspects:

Basic Logic

A basic version of the logic which is compatible with the meaning representation chosen for documents and queries should be produced. For instance, if our model for representing documents contemplates a representation in terms of sets of keywords, it will be useful to investigate the fragment of the chosen logic that deals with implications whose antecedents may only be in the form of a conjunction of simple propositions. This logic should be amenable to extensions, as indicated in the following points, and it is therefore essential its being equipped with a formal and intuitive semantics. This is the case of $\mathbf{E_{fde}}$, whose four-valued semantics was developed independently by Belnap [2,3] and Dunn [4].

Computational Properties

The computational properties of the logic should be investigated. As the logic is the theoretical basis of a software system which must deal with on-line requests, it must allow efficient implementations of its inference apparatus.

This investigation should produce a mapping, taking classes of the decision problem of the logic ("Does α entail β?") which are significant to Information Retrieval, into computational classes, ranging from the undecidable to complexity classes [7]. This mapping will play a central role in the design of the inference engine of the Information Retrieval system: should intractability arise for relevant classes of the Information Retrieval problem, *ad hoc* heuristics will have to be devised to cope with such cases.
The computational properties of $\mathbf{E_{fde}}$ have been partially investigated. While deciding entailment in the general case is likely to be intractable (technically, the problem is co-NP-complete [11]), whenever α and β are formulae in Conjunctive Normal Form, there exists a $O(|\alpha| \cdot |\beta|)$ algorithm that tests if $\alpha \rightarrow \beta$ holds [9]. While this provides further evidence to the suitability of $\mathbf{E_{fde}}$ as a logic for Information Retrieval, the investigation on the complexity of tautological entailment must clearly be extended to other classes of formulae. Turning to quantification, first-order tautological entailment is, unfortunately, powerful enough so that first-order logical implication can be reduced to it, thus sanctioning its undecidability [11]. There are, however, neighbor notions of entailment that are decidable, and that have been used in Knowledge Representation to model tractable forms of reasoning on beliefs [8]. The suitability of these notions to Information Retrieval must be investigated, pursuing further the study of their computational properties, in case they turn out to be appropriate for Information Retrieval.

Extensions

The basic logic should be extended with notions that are significant to Information Retrieval. While the decision problem we have mentioned in the previous point takes the form of a *yes/no* question, it is often the case, in designing an Information Retrieval system, that one wants to measure how well a document fits a query. On the basis of this measure, documents that do not qualify but that are "not too far" from the query may be accepted in the query result, whereas documents that qualify but "worse" than others may be rejected. Such a notion has to be embedded in the logic in a principled form, otherwise the meaning of inferences will be lost. As anticipated in the first point above, the logic $\mathbf{E_{fde}}$ lends itself to extensions very well, as it has a clean and intuitive semantics, which can be used as a basis of any further addition to the world of relevance.

Multiple representations

The relevance logic-based representation and reasoning should be integrated with other representations of documents, to which other inference mechanisms are associated. It is believed that the current limitations of Information Retrieval systems may be partly overtaken by providing such systems with multiple representations of the same document, which allow the exploitation of the different perceptions different users typically have of the same document, or the same user has of different documents. As an example, the system may support the representation of a document as a structured, complex object, whose components are objects themselves, which, being meaningful to the users, can be used in stating the users' requests. Furthermore, the user may have some knowledge about a document and its contents that he/she wants the system to be able to store and (in a limited form) reason about. This is especially the case with multimedia components such as images, where an *interpretation* is usually associated to an image so that the system can answer to contents-oriented requests on that image. For all these representations, a standard logical representation seems to be appropriate, as it has emerged from the conceptual modelling of database applications [12]. So, in fact, what this aspect amounts to is the inclusion in the logic being discussed of a fragment of first-order logic, which be expressively adequate to the alternative representations of a document, and computationally tractable with respect to the kind of inferences required on these representations.

4.3 Engineering

An *implementative* study should then be undertaken, aimed at the construction of IR systems that comply with the logic produced in preceding step, and at the subsequent empirical testing in real IR contexts.

References

[1] Alan R. Anderson and Nuel D. Belnap. *Entailment - the logic of relevance and necessity.* Princeton University Press, Princeton, NJ, 1975.

[2] Nuel D. Belnap. How a computer should think. In *Contemporary Aspects of Philosophy: Proceedings of the Oxford International Symposium*, pages 30–56, Oxford, GB, 1975.

[3] Nuel D. Belnap. A useful four-valued logic. In G. Epstein and J. M. Dunn, editors, *Modern Uses of Multiple-Valued Logic*, pages 8–37, Reidel, Dordrecht, NL, 1977.

[4] J. Michael Dunn. Intuitive semantics for first-degree entailments and coupled trees. *Philosophical studies*, 29:149–168, 1976.

[5] J. Michael Dunn. Relevance logic and entailment. In Dov Gabbay and Franz Guenthner, editors, *Handbook of Philosophical Logic. Vol. III*, pages 117–224, Reidel, Dordrecht, NL, 1986.

[6] Susan Haack. *Philosophy of logics.* Cambridge University Press, Cambridge, GB, 1977.

[7] David Johnson. A catalog of complexity classes. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science. Vol. A Algorithms and Complexity*, pages 67–161, The MIT Press/Elsevier, Cambridge, MA, 1990.

[8] Gerhard Lakemeyer. Step towards a first-order logic of explicit and implicit belief. In *Proceedings of the First Conference on Theoretical Aspects of Reasoning About Knowledge*, pages 325–340, Asilomar, CA, 1986.

[9] Hector J. Levesque. A logic of implicit and explicit belief. In *Proceedings of AAAI-84, 4th Conference of the American Association for Artificial Intelligence*, pages 198–202, Austin, TX, 1984.

[10] E.J. Nelson. On three logical principles in intension. *The Monist*, 43, 1933.

[11] Peter F. Patel-Schneider. *Decidable, Logic-Based Knowledge Representation.* Technical Report 201/87, Department of Computer Science, University of Toronto, Toronto, Ontario, 1987.

[12] Raymond Reiter. Towards a logical reconstruction of relational database theory. In Michael L. Brodie, John Mylopoulos, and Joachim W. Schmidt, editors, *On Conceptual Modelling*, pages 191–233, Springer Verlag, New York, NY, 1984.

[13] C. J. van Rijsbergen. Towards an information logic. In *Proceedings of SIGIR-89, 12th Conference of the ACM Special Interest Group on Information Retrieval*, pages 77–86, Cambridge, MA, 1989.