Incremental Knowledge Acquisition for Non-Monotonic Reasoning

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

The use of conventional non-monotonic reasoning tools in real-sized knowledge-based applications is hindered by the fact that the knowledge acquisition phase cannot be accomplished in the incremental way that is instead typical of knowledge base management systems based on monotonic logics. As a result, some researchers have departed from orthodox non-monotonic formalisms and proposed languages for the representation of *Multiple Inheritance Networks with Exceptions* (MINEs). Such languages do not suffer from the problem of incrementality in knowledge acquisition, but are inadequate both from a formal and from an empirical point of view. In fact, they are not endowed with a formal semantics, and the intuitions that underlie their inferential mechanisms are far from being widely agreed upon.

In this paper we discuss an approach to non-monotonic reasoning which does allow the phase of knowledge acquisition to be accomplished in an incremental and modular way, but at the same time relies on a solid and widely acknowledged formal apparatus such as First Order Logic (FOL). We have obtained this by specifying a (non-monotonic) function that maps MINEs into sets of FOL formulae. We have shown that the mapping function we discuss is sound and complete, in the sense that each conclusion that can be derived from a MINE is also derivable from the set of FOL formulae resulting from its translation via the mapping function, and vice-versa.

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

1 Introduction

Traditional formal systems have the property of being *monotonic*, in the sense that an addition to the set of formulae making up the knowledge base (KB) never determines a decrement in the set of conclusions that may be derived from the KB via the inference rules. Formally, if $\Gamma \vdash A$ holds, then necessarily $\Gamma \cup \Gamma' \vdash A$ holds, where Γ and Γ' are sets of formulae, A is a formula and \vdash is the derivability relation.

Unfortunately, only few application domains exist whose formalization may be accomplished by means of monotonic formal systems¹. In fact, there are many reasoning patterns for which a monotonic formal system would prove inadequate. One of these is *default reasoning*, a reasoning style that is to be applied whenever the rules involved allow for *exceptions*. Such rules dictate that, "by default" (i.e. whenever there is no information to the contrary), one may assume that the situation that is to be dealt with is a *typical* (i.e. not exceptional) situation. However, a possible subsequent acquisition of new knowledge might lead one to discover that this situation is not typical as was believed, but *exceptional*; in this case, what had been assumed "by default" is *revoked*. Hence, reasoning by default (also known as *reasoning in the presence of exceptions*) is a non-monotonic kind of reasoning.

An instance of this reasoning style is the inferential chain induced by the assertion "if x is a bird, assume that it flies, unless you know it does not". If we reason according to this rule and know that Opus is a bird, in the absence of information to the contrary we may conclude that Opus flies; but if we later learn that Opus is a penguin, knowing that penguins typically do not fly, we have to revoke the conclusion that Opus flies.

Non-monotonic reasoning is nowadays a hot topic in AI; in fact, the spectrum of its applications is extremely wide, ranging from image interpretation [11], reasoning about action [10], troubleshooting [13], knowledge bases automatic generation [15] and natural language understanding [2]. The requirements that are imposed by these application domains have lead to the development of a variety of non-monotonic formalisms, most of which belong to the offspring of Doyle and McDermott's Nonmonotonic Logic, Reiter's Default Logic and McCarthy's Circumscription². Unfortunately, these formalisms, besides having unattractive metalogical and computational properties (i.e. undecidability of the versions with quantifiers, computational intractability of those without quantifiers), suffer from a problem that hinders their use in those knowledge representation contexts which require that the knowledge acquisition phase be accomplished in an incremental and modular fashion. We will call this problem the *Exceptions Explicitation Problem* (EEP).

Incrementality of knowledge acquisition is an asset of knowledge base management systems that hardly needs to be argued for. Large knowledge bases are the result of an evolutionary process; this happens because knowledge entry is a time-consuming process, and because knowledge may simply become available at later stages of the process, possibly contradicting (or specializing) previously acquired knowledge. When a large KB is built

¹Witness the fact that most research papers about monotonic KR languages draw their examples from "unrealistically neat" application domains, such as e.g. geometry, kinships, and the like.

²Up-to-date accounts of what is going on in the field may be gathered by consulting [1, 4, 12].

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 a possibly ensuing deletion) of pre-existing chunks; in other words, it is desirable that knowledge acquisition be incremental (or *compositional*).

In Section 2 of this paper we will discuss how the EEP manifests itself in standard non-monotonic formalisms, and will see how this makes their implementation into working KBMSs allowing for incremental knowledge acquisition virtually impossible.

In Section 3 we will briefly describe a class of languages, those allowing the representation of *Multiple Inheritance Networks with Exceptions* (MINEs – see e.g. [16, 19]), that solve this problem by a technique we will call *implicit handling of exceptions*. Unfortunately, such languages lack a denotational semantics that account for the meaning of the linguistic primitives contained therein in a clear and unambiguous way. Besides, given the complexity and awkwardness of the non-standard proof theory such languages rely on in order to specify the conclusions that are derivable from a KB, it is substantially implausible to think that such a proof theory might "implicitly" constitute a semantics for the languages in question³.

In order to overcome these deficiencies, in Section 4 we will propose an approach to non-monotonic reasoning that combines the advantages (in terms of incrementality of the knowledge acquisition phase) that are offered by the languages for MINEs, with those (in terms of semantic clarity) that are offered by a formally solid apparatus such as First Order Logic (FOL). This approach is accomplished by specifying a (non-monotonic) function that maps a MINE into a set of FOL formulae. We will show that this mapping function is a "good" mapping, in the sense that each conclusion which is derivable from a MINE is also derivable from the set of FOL formulae that results from the application of the mapping function to the MINE, and vice-versa. An interesting side-effect of this result is that MINE-like reasoning can be performed by means of ordinary, ready-made first order theorem provers, with no need to add specific machinery. Section 5 concludes.

2 The Exceptions Explicitation Problem

In this section we will discuss, by means of a concrete example, how the EEP manifests itself in the context of Nonmonotonic Logic (NML); to this respect, other formalisms such as Default Logic and Circumscription behave in a completely analogous way, and will not be discussed. The "penguin" example that we have hinted at in the preceding section may be represented by means of the NML formula

$$\forall x \; Bird(x) \land M[Flies(x)] \Rightarrow Flies(x) \tag{1}$$

(where $M[\alpha]$ informally means " α does not contradict the KB"). Let us consider an

 $^{^{3}}$ For a more detailed discussion on this "implicit" approach to the semantics of representation languages see e.g. [3].

exception to the rule, "penguins are birds that typically do not fly", and let us represent it by means of the axioms:

$$\forall x \ Penguin(x) \Rightarrow Bird(x) \tag{2}$$

$$\forall x \ Penguin(x) \land M[\neg Flies(x)] \Rightarrow \neg Flies(x) \tag{3}$$

There is a problem hidden in this set of formulae: in NML the addition of the assertion Penguin(opus) to axioms (1)-(3) generates two different "sets of conclusions" ("fixpoints", in NML terminology): in one of them Opus flies while in the other Opus does not fly, depending on whether we "first" consider axiom (1) or axiom (3), respectively; this happens regardless of the fact that our knowledge of the animal world makes us strongly prefer the conclusion according to which Opus does not fly.

Hence NML seems, at first sight, inadequate to account for this kind of reasoning, as it generates a situation of ambiguity in the representation of a state of affairs that is all but ambiguous to us^4 . At first sight, it seems possible to overcome this problem without departing from NML; what one needs to do is *explicitly* consider exceptions in each default formula they involve. Accordingly, a set of axioms that induces a correct behaviour of NML for our example is the following:

$$\forall x \; Bird(x) \land M[Flies(x) \land \neg Penguin(x)] \Rightarrow Flies(x)$$

$$\forall x \; Penguin(x) \Rightarrow Bird(x)$$

$$\forall x \; Penguin(x) \land M[\neg Flies(x)] \Rightarrow \neg Flies(x)$$

$$(4)$$

At this point, given Penguin(opus), axiom (4) does not allow one to conclude Flies(opus); this leaves us with one fixpoint only, a fixpoint that contains only the conclusions we actually subscribe to.

However, this solution looks like an *ad hoc* patch rather than a real solution. In fact, it should be noted that, given the informal meaning that Doyle and McDermott attribute to NML formulae, the conclusion according to which Opus does not fly should already follow from axioms (1)-(3), as these already represent the fact that penguins are exceptional birds as far as the ability to fly is involved; consequently, the modification of axiom (1) to yield (4) seems a redundant re-assertion of this fact.

Besides, one should note that, until the KB did not contain any reference to penguins, axiom (1) was more than adequate. The introduction of axioms (2) and (3) has obliged us to transform axiom (1) into axiom (4); in the more general case, the acquisition of new knowledge obliges us to *revoke* (in order to subsequently assert some modified version of them) *previously acquired formulae from the KB*. It is then evident that the realization of default reasoning by means of NML is undermined by the impossibility to update KBs by

 $^{^{4}}$ The generation of multiple fixpoints in NML may otherwise be due to an *inherent* ambiguity of the state of affairs being represented. A very famous example of this is the so-called "Nixon diamond" (see Example 2).

simple piecemeal additions of new knowledge chunks. In semantic terms this means the impossibility, unlike what happens with traditional logical formalisms, to characterize the extensional semantics of the modified KB in a compositional way, i.e. as a combination of the extensions of the original KB and of the newly acquired formulae.

But the need to explicitly represent exceptions brings about other, even more compelling problems, having to do with the way in which, at the time of the introduction of a new formula into the KB, it is to be determined which are the formulae that are to be modified and which are the ones that are not. It turns out that this operation cannot be realized by simply performing a number of matchings between the formula to be introduced and the other formulae in the KB; on the contrary, it requires in general a NML theorem proving *operation.* This may clearly be seen by taking a look at our example: the relation between the precondition of axiom (1) and the precondition of axiom (3) has been determined by relying on the fact that Bird(x) is derivable from Penquin(x) by means of axiom (2); the transformation of axiom (1) into axiom (4) is then realized in order to "inhibit" not axiom (1) in itself, but the inferential chain that would lead us to conclude Flies(x)from Penquin(x). In general, given a NML knowledge base, it is not even determined a priori whether there exist one, several or no inferential chains that lead to the undesired conclusion (in our case, one, several or no inferential chains between Penguin(x) and Flies(x); in case several such chains exist, each of them has to be inhibited by means of a call to the NML theorem prover, each resulting in the introduction of a subformula of type $M[\alpha]$.

That the phase of KB construction requires repeated calls to the NML theorem prover in order to maintain the consistency of the KB that is being built, is in itself rather implausible; but the whole endeavour becomes absurd once we consider that NML is undecidable! This means that, unless the construction of the KB is realized in a completely *static* (non incremental) way, the problem of knowledge acquisition in NML (and, analogously, in the other non-monotonic formalisms mentioned in Section 1) is an *unsolvable* problem.

3 Multiple Inheritance Networks with Exceptions

The Exceptions Explicitation Problem that affects standard non-monotonic formalisms has lead to the definition of languages for the representation of *Multiple Inheritance Networks* with Exceptions (MINEs). Such languages have two fundamental characteristics:

- they are specifically oriented to the representation of (non-monotonic) knowledge of a *taxonomic* nature (i.e. of knowledge that lends itself to being organized hierarchically); hence, they are less expressive of those we have mentioned in Section 1, in the sense that they only allow for monadic predicate symbols, a limited use of negation, and no disjunction at all;
- they implement an *implicit* handling of exceptions (see below).

A MINE is a directed graph whose nodes represent classes or individuals of the domain of discourse, and whose edges represent relationships of "conceptual containment" between

the nodes involved. For instance, a MINE Γ might contain the nodes *opus*, *Penguin* and *Bird*, and contain the edges *opus* \rightarrow *Penguin*, *Bird* \rightarrow *Flies* and *Penguin* $\not\rightarrow$ *Flies*. The word *inheritance* refers to the fact that the deductive mechanisms of these languages are such that an edge *Bird* \rightarrow *Flies* has the effect of transmitting "by inheritance" all the properties of *Flies* to *Bird*; this inheritance is *multiple* because a MINE Γ may contain more than one edge outgoing from the same node, the net effect being that the node inherits properties from different, unrelated nodes. Last, we speak of networks *with exceptions* because the intended denotation of an edge such as $p \rightarrow q$ is the set of states of affairs in which "p's are typically q's"; this assertion allows for the existence of p's that are not q's, i.e. of exceptions. Dual arguments apply to edges of type $p \not\rightarrow q$.

The languages for the representation of MINEs that have been proposed in the literature are by now a fairly vast and multifaceted class (see e.g. [5, 6, 14, 18]). It is not our purpose to give an exhaustive review of the constructs and inferential mechanisms that appear in the various languages (to this respect, the interested reader may consult [16, 19]); what we want is to introduce a number of fundamental concepts that, besides being common to a vast class of MINE representation languages, will be sufficient to illustrate the problems affecting such languages in general, and to motivate the approach described in Section 4.

In the discussion that follows we will restrict ourselves to the case in which the relation $\rightarrow \cup \not\rightarrow ($ a relation which we will indicate with the symbol $\sim)$ is irreflexive and transitive (i.e. a *strict partial order*); this limitation is not particularly stringent from a pragmatic point of view. In MINEs that enjoy this property, exceptions are *implicitly* handled by this ordering: to a first approximation, we can say that, in case of conflicts, an edge is "preferred" to another if the source node of the first precedes the source node of the second in the ordering. For example, given the MINE consisting of the edges *opus* \rightarrow *Penguin*, *Penguin* \rightarrow *Bird*, *Bird* \rightarrow *Flies*, *Penguin* $\not\rightarrow$ *Flies*, we will see how the conclusion according to which Opus does not fly is given priority, on the grounds that the premise of *Penguin* $\not\rightarrow$ *Flies* is more specific than that of *Bird* \rightarrow *Flies*, and that, as a consequence, the conclusion derivable from the former is more reliable than the one derivable from the latter⁵.

Let us now formally define our language for the representation of MINEs.

Definition 1 Let \mathcal{I} and \mathcal{C} be two finite sets of individual constants and predicate symbols, respectively, and let \mathcal{A} be a finite set of positive edges $x \to y$ and negative edges $x \neq y$, where $x \in \mathcal{I} \cup \mathcal{C}$ and $y \in \mathcal{C}$. A Multiple Inheritance Network with Exceptions (MINE) is a triple $\Gamma = \langle \mathcal{I}, \mathcal{C}, \mathcal{A} \rangle$ such that the relation $\to \cup \neq$ is a strict partial order. The elements of the set $\mathcal{I} \cup \mathcal{C}$ are called nodes. The polarity of an edge is given by its positive or negative sign.

⁵The reader may have noticed that in this example we have used the edge $Penguin \rightarrow Bird$ to formalize a fact that would be better represented by an implication of classical logic, given that *all* penguins are birds. The integration between "categorical" and "default" implications in a single framework is not among the purposes of this paper; for simplicity, we will hereafter take all implications as having a default nature. For a more detailed discussion on the above-mentioned integration see e.g. [5].

With Γ_I , Γ_C and Γ_A we will refer to the components \mathcal{I} , \mathcal{C} and \mathcal{A} of Γ . In the examples that follow we will use lowercased words as metavariables ranging over individual constants, capitalized words for predicate symbols and the letters p, q, \ldots, x, y, z for (generic) nodes.

Let us now define, given a MINE Γ , what the set of conclusions that may be derived from Γ is; this is very important because, similarly to what happens in other logics for knowledge representation, it establishes the way in which we may infer knowledge *implicit* in the KB from the knowledge explicitly coded therein. However, in order to do so we have to introduce the notions of "chain" and "path".

Definition 2 A chain in a MINE Γ is a sequence of edges in Γ_A , $x_1 \rightsquigarrow x_2 \rightsquigarrow \ldots \rightsquigarrow x_{n-1} \rightsquigarrow x_n \ (n \ge 2)$. The node x_1 is called initial node of the chain and x_n is called terminal node. The degree of a chain in Γ with initial node x and terminal node y is given by the length of the longest chain from x to y in Γ .

Hereafter we will use lowercased greek letters as metavariables ranging over chains, and we will indicate with $deg_{\Gamma}(\sigma)$ the degree of a chain σ on Γ .

Definition 3 A positive path in a MINE Γ is a chain $x_1 \to x_2 \to \ldots \to x_{n-1} \to x_n$ $(n \geq 2)$ in Γ consisting of positive edges only. A negative path in a MINE Γ is a chain $x_1 \to x_2 \to \ldots \to x_{n-1} \not\to x_n$ $(n \geq 2)$ in Γ consisting of a (possibly empty) sequence of positive edges followed by one negative edge. The polarity of a path is given by the polarity of its last edge. The edge formed by the initial node of a path, its terminal node and its polarity is called conclusion derived from the path.

Intuitively, paths are those chains from which it is somehow possible to draw conclusions. For example, let us consider a MINE such as $KB1 = \langle \{opus\}, \{Penguin, Bird, Flies, Swims\}, \{opus \rightarrow Penguin, Penguin \rightarrow Bird, Bird \rightarrow Flies, Penguin \not\rightarrow Flies, Flies \not\rightarrow Swims\}\rangle$. The chain

$$opus \to Penguin \to Bird \to Flies$$
 (5)

is a positive path, from which we can draw the conclusion $opus \rightarrow Flies$ ("As there is no information to the contrary, we assume that Opus flies"), while from the negative path

$$opus \to Penguin \not\to Flies$$
 (6)

we may conclude $opus \not\rightarrow Flies$ ("As there is no information to the contrary, we assume that Opus does not fly"). The chain

$$opus \rightarrow Penguin \not\rightarrow Flies \not\rightarrow Swims$$
 (7)

is not a path; in fact, no reasonable conclusion may be drawn from it with respect to the fact whether Opus can swim or not. Similar empirical considerations have induced researchers to define the notion of "conclusion" with respect to paths only, and not with respect to chains in general.

As we have seen, MINEs such as KB1 have paths from which contradictory conclusions may be derived; unfortunately, this generates situations of ambiguity in the representation of states of affairs which are all but ambiguous to us. In order to eliminate these situations, what we will do is to define which are, given a MINE Γ , the paths from which we may derive "reliable" conclusions, and that must consequently be given priority over others. For example, in a MINE such as KB1 this will allow us to give path (6) priority over path (5), hence to inhibit the undesired conclusion according to which Opus flies. In MINE terminology, paths from which we may derive "reliable" conclusions are called *derivable* paths.

Definition 4 A positive path σ is derivable from a MINE Γ (in symbols: $\Gamma \vdash \sigma$) if and only if the following conditions obtain:

- 1. if $\sigma = x \rightarrow y$, then $\Gamma \vdash \sigma$ if and only if $\sigma \in \Gamma_A$ and $x \not\rightarrow y \notin \Gamma_A$;
- 2. if $deg_{\Gamma}(\sigma) = n > 1$, then $\Gamma \vdash \sigma$ if and only if there exists a (possibly empty) positive path δ and nodes x, y and $u \in \Gamma_I \cup \Gamma_C$ such that:
 - (a) $\sigma = x \to \delta \to u \to y$
 - (b) $u \to y \in \Gamma_A$;
 - (c) $\Gamma \vdash x \to \delta \to u;$
 - (d) $x \not\rightarrow y \notin \Gamma_A$;
 - (e) for all v and for all τ such that $\Gamma \vdash x \to \tau \to v$ and $v \not\to y \in \Gamma_A$, there exists a z such that $z \to y \in \Gamma_A$ and either z = x or $\Gamma \vdash x \to \gamma_1 \to z \to \gamma_2 \to v$ for some path γ_1 and for some path γ_2 .

Definition 5 A negative path σ is derivable from a MINE Γ (in symbols: $\Gamma \vdash \sigma$) if and only if the following conditions obtain:

- 1. if $\sigma = x \not\rightarrow y$, then $\Gamma \vdash \sigma$ if and only if $\sigma \in \Gamma_A$ and $x \rightarrow y \notin \Gamma_A$
- 2. if and only if $deg_{\Gamma}(\sigma) = n > 1$, then $\Gamma \vdash \sigma$ if and only if there exists a (possibly empty) positive path δ and nodes x, y and $u \in \Gamma_I \cup \Gamma_C$ such that:
 - (a) $\sigma = x \to \delta \to u \not\to y$
 - (b) $u \not\rightarrow y \in \Gamma_A$
 - (c) $\Gamma \vdash x \to \delta \to u$
 - (d) $x \to y \notin \Gamma_A$
 - (e) for all v and for all τ such that $\Gamma \vdash x \to \tau \to v$ and $v \to y \in \Gamma_A$ there exists a z such that $z \not\to y \in \Gamma_A$ and either z = x or $\Gamma \vdash x \to \gamma_1 \to z \to \gamma_2 \to v$ for some path γ_1 and for some path γ_2 .

Following the terminology of [19] we might say that clauses (b) and (c) of Definitions 4 and 5 represent the option for a style of reasoning informed by "bottom-up chaining", while clause (e) represents the option for "off-path preclusion"; the mapping function that we will describe in Section 4 subscribes to these options⁶.

Definition 6 The set of conclusions of a MINE Γ (in symbols: $C(\Gamma)$) is the set of those edges $p \rightsquigarrow q$ such that $p \rightsquigarrow q$ is the conclusion of a path σ derivable from Γ .

The set of conclusions that may be drawn from a MINE Γ is the set that represents the knowledge "implicitly" present in the KB.

4 Mapping MINEs into FOL

The implicit handling of exceptions implemented in MINE representation languages by means of the ordering of nodes in the network allows the phase of knowledge acquisition to be realized in an incremental way: recalling the example of Section 2, if our KB contained the edge $Bird \rightarrow Flies$, a subsequent introduction of the edges $opus \rightarrow Penguin$, $Penguin \rightarrow Bird$ and $Penguin \not\rightarrow Flies$ would not bring about the need to modify the pre-existing edge, as the new strict partial order deriving from the introduction of the new edges would inhibit the undesired conclusion $opus \rightarrow Flies$.

The limit of the MINE-based approach is its lack of semantic clarity. In fact, the development of a model-theoretic semantics for MINE representation languages is still an open problem, and the lack of such a semantics makes both the analysis of these languages and the comparison between them extremely difficult.

In this paper we discuss an approach to non-monotonic reasoning that allows for incremental knowledge acquisition, while at the same time relying on a formally solid and widely acknowledged framework such as First Order Logic. We have accomplished this by specifying a first order representation of the MINE formalism (a representation that we dub LT, the Logical Theory of Multiple Inheritance with Exceptions), and a mapping function \mathcal{M} that maps a MINE Γ into a set of FOL formulae. The LT theory will be on all counts our formalism for default reasoning; a mathematical theory for default reasoning (i.e. a KB) will be obtained by adding to LT the result of the application of \mathcal{M} to a particular MINE, obtaining the FOL knowledge base $\mathcal{T}(\Gamma) = \mathcal{M}(\Gamma) \cup LT$. The net effect is that the LT component of $\mathcal{T}(\Gamma)$ remains fixed throughout the phase of knowledge acquisition (in the same sense in which the logical axioms are fixed for a given logical system), while the only component that varies is the $\mathcal{M}(\Gamma)$ component.

⁶The mapping function may be modified without difficulty in case one wants to opt for top-down chaining and/or on-path preclusion. Our notion of derivability is very similar to the one proposed in [6]; it differs from it in the fact that in [6] item (1) of the two definitions specifies that, if $deg_{\Gamma}(\sigma) = 1$, then $\Gamma \vdash \sigma$ if and only if $\sigma \in \Gamma_A$ (i.e. unlike in our definition, it is not necessary that the edge with opposite polarity does not belong to Γ_A). The modifications that could be made to our formalism in order to obtain the behaviour described in [6] would be trivial, and will not be considered.

It may be shown that \mathcal{T} is sound and complete, in the sense that each conclusion derivable from a MINE Γ is also derivable from $\mathcal{T}(\Gamma)$, and vice-versa.

At this point a question arises naturally: how can a non-monotonic system be mapped into FOL, a notoriously monotonic system, while at the same time maintaining nonmonotonicity? Isn't this a contradiction in terms? In order to hint that this is a plausible state of affairs, let us recall that Circumscription (one of the most credited non-monotonic formalisms) is nothing else but a mapping function from theories of a formalism F_1 (FOL) into theories of another formalism F_2 (second order logic), and that in this case too the target formalism is monotonic. In our case (and in the case of Circumscription too), nonmonotonicity is a property of the mapping function, and not of the target formalism. In fact, if Γ and Γ' are MINEs such that $\Gamma \cup \Gamma'$ is the MINE obtained through the union of the respective components (that is: $\Gamma \cup \Gamma' = \langle \Gamma_I \cup \Gamma'_I, \Gamma_C \cup \Gamma'_C, \Gamma_A \cup \Gamma'_A \rangle$), \mathcal{M} is such that the translation of Γ is not in general a logical consequence of the translation of $\Gamma \cup \Gamma'$; that is, in general $\mathcal{M}(\Gamma \cup \Gamma') \not\models \mathcal{M}(\Gamma)$. This is a necessary condition for any function aspiring to translate a non-monotonic formal system into a monotonic one; if S_1 is a non-monotonic system based on language \mathcal{L} , E and E' are sets of expressions on \mathcal{L} , and \mathcal{K} is the translation function of sets of expressions of \mathcal{L} into sets of expressions of a monotonic system \mathcal{S}_2 , then \mathcal{K} must be such that in general it is not the case that $\mathcal{K}[E \cup E'] \models_{S_2} \mathcal{K}[E]$.

In order to prove our result, we will proceed in the following way. First of all, we will define the mapping function \mathcal{M} ; we will then define LT, the Logical Theory of Multiple Inheritance with Exceptions, and will then show that belonging to the set of conclusions of a MINE Γ is equivalent to being a logical consequence in LT (a notion we will indicate with the symbol " \models_{LT} ") of $\mathcal{M}\Gamma^7$. All this will be described in a semi-informal way in Sections 4.1 and 4.2, and subsequently formalized in Section 4.3.

4.1 The \mathcal{M} function

The purpose of the \mathcal{M} function is that of translating a MINE $\Gamma = \langle \mathcal{I}, \mathcal{C}, \mathcal{A} \rangle$ into a set of formulae $\mathcal{M}(\Gamma)$ constituting a first order representation of the graph formed by the nodes and edges of Γ . In order to obtain this, first of all let us represent every node pin $\Gamma_I \cup \Gamma_C$ by an individual constant p of the first order language. The edges of Γ_A are instead represented through instances of the binary predicates ISDA ("IS Directly A") and ISDNA ("IS Directly Not A"): the formulae ISDA(p,q) and ISDNA(p,q) represent then the fact that edges $p \to q$ and $p \neq q$, respectively, belong to Γ_A . Let us notice that, in order for Γ_A to be represented adequately, the following conditions must obtain:

- the predicate *ISDA* must be such that $p \to q \in \Gamma_A$ if and only if every model of $\mathcal{M}(\Gamma)$ is also a model of ISDA(p,q);
- the predicate ISDNA must be such that $p \not\rightarrow q \in \Gamma_A$ if and only if every model of $\mathcal{M}(\Gamma)$ is also a model of ISDNA(p,q).

⁷Let us recall that in the language of mathematical logic, $\alpha \models_A \beta$ is short for $A \cup \{\alpha\} \models \beta$, where α and β are formulae and A is a theory of the language in question.

In order to accomplish this, let us consider for every edge $p \to q \in \Gamma_A$ a formula of type ISDA(p,q). It will be apparent in the proof of Lemma 1 that we need these formulae to be all and *the only* instances of the ISDA predicate to be true in the models of $\mathcal{M}(\Gamma)$; in order to accomplish this, we impose that the following axiom holds⁸:

$$\forall x \forall y \ ISDA(x, y) \Leftrightarrow \bigvee_{p \ \rightarrow q \in \Gamma_A} (x = p \land y = q)$$
(8)

Analogously, if we want to represent negative edges by means of the predicate ISDNA, we obtain the formula

$$\forall x \forall y \; ISDNA(x, y) \Leftrightarrow \bigvee_{p \neq q \in \Gamma_A} (x = p \land y = q) \tag{9}$$

In order to obtain a one-to-one correspondence between the network and its FOL representation, we have also to consider that, in any MINE, nodes are considered as symbols referring to individuals "implicitly" different from each other; this is often an unstated assumption in such languages, but it is nevertheless "wired" in the reasoning machinery that implements them. As a consequence, also individual constants must be interpreted as referring each to a different individual: that is, we do not want that models of $\mathcal{M}(\Gamma)$ may exist that associate syntactically different constants to the same individual of the domain of discourse. In order for this to hold, we impose that the axiom

$$\bigwedge_{\{p,q\}\in\Gamma_I\cup\Gamma_C} (p\neq q) \tag{10}$$

holds, an axiom that formalizes what is usually known as the "unique names assumption".

This concludes the definition of the \mathcal{M} function, a function that may then be concisely described by means of the following definition.

Definition 7 Let Γ be a MINE; \mathcal{M} is the function that maps Γ into the FOL axiom

$$(\forall x \forall y \ ISDA(x, y) \Leftrightarrow \bigvee_{p \to q \in \Gamma_A} (x = p \land y = q)) \land (\forall x \forall y \ ISDNA(x, y) \Leftrightarrow \bigvee_{p \neq q \in \Gamma_A} (x = p \land y = q)) \land (\land_{\{p,q\} \in \Gamma_I \cup \Gamma_C} (p \neq q))$$

Note that the size of the axiom is at most $O(n^2)$, where n is the number of nodes in Γ .

At this point we may already discuss how incrementality is achieved in our framework. In the preceding section we have hinted at the fact that the LT component of $\mathcal{T}(\Gamma)$ remains

⁸This axiom turns out to be equivalent [9] to the minimization of the predicate ISDA by means of the Circumscription axiom $Circum(\bigwedge_{p\to q\in\Gamma_A} ISDA(p,q); ISDA; ())$; alternatively, this minimization may be seen as the completion of the ISDA predicate.

fixed throughout the phase of knowledge acquisition, while the only component that varies is the $\mathcal{M}(\Gamma)$ component. We may now see that the way in which this component varies is informed by the principles of incrementality and compositionality. In fact, if one needs to add the equivalent of an edge $r \to s$ (resp. $r \neq s$) to a knowledge base $\mathcal{T}(\Gamma)$, one needs only to add the formula $(x = r \land y = s)$ as a further operand of the disjunction $\bigvee_{p\to q\in\Gamma_A}(x = p \land y = q)$ (resp. $\bigvee_{p\neq q\in\Gamma_A}(x = p \land y = q)$), and the formula $(r \neq s)$ to the conjunction $\bigwedge_{\{p,q\}\in\Gamma_I\cup\Gamma_C}(p\neq q)$. As required, knowledge acquisition becomes a matter of simple piecemeal additions of new chunks of knowledge to the pre-existing KB.

4.2 The Logical Theory of Multiple Inheritance with Exceptions

Now that we have defined the \mathcal{M} function, let us build a first order theory LT such that a formula of type ISA(x, y) (resp. ISNA(x, y)) is derivable from $\mathcal{M}(\Gamma)$ in LT if and only if $x \to y$ (resp. $x \neq y$) belongs to $C(\Gamma)$, the set of conclusions of Γ . The LT theory will roughly consist of the definitions of the predicates ISA and ISNA in terms of the predicates ISDA and ISDNA we have seen in the preceding section. This definition is rather complex; hence, it will be introduced in an incremental way, stepping through the definition of some other auxiliary predicates representing the existence of particular conditions on the topology of Γ .

In order to do so, let us first start by defining two predicates $\Pi_P(x, t, y)$ and $\Pi_N(x, t, y)$ representing the derivability in Γ of a path (positive or negative, respectively) with initial node x and terminal node y passing from node t. As a consequence, the following conditions must obtain:

- the predicate Π_P must be such that $\Gamma \vdash x \to \gamma_1 \to t \to \gamma_2 \to y$, for some γ_1 and γ_2 , if and only if every model of $\mathcal{M}(\Gamma) \cup LT$ is also a model of $\Pi_P(x, t, y)$;
- the predicate Π_N must be such that $\Gamma \vdash x \to \gamma_1 \to t \to \gamma_2 \not\to y$, for some γ_1 and γ_2 , if and only if every model of $\mathcal{M}(\Gamma) \cup LT$ is also a model of $\Pi_N(x, t, y)$.

In order for Π_P and Π_N to have the desired characteristics, let us define them by means of a "literal" translation of Definitions 4 and 5. The translation of Definition 4 is obtained by imposing that the following formula (which has been subdivided in several numbered subformulae and indented so as to highlight its structural affinity with the definition) belongs to LT.

$$\forall x \forall t \forall y \ \Pi_P(x,t,y) \Leftrightarrow \\ ((x = t \land ISDA(x,y) \land \neg ISDNA(x,y)) & 1. \\ \lor \ (((\exists t'\Pi_P(x,t,t') \land t \neq t' \land ISDA(t',y)) & 2.(a) - (b) \\ \lor \ (\exists t'\Pi_P(x,t',t) \land t \neq t' \land ISDA(t,y))) \land \\ \neg ISDNA(x,y) & (d) \\ \land \ (\forall v \exists v'(\Pi_P(x,v',v) \land ISDNA(v,y)) & (e) \\ \Rightarrow (\exists z ISDA(z,y) \land (z = x \lor \Pi_P(x,z,v))))))$$

At this point we may characterize the conclusion that can be drawn from a positive path, by means of a predicate ISA(x, y), by adding to LT the simple formula

$$\forall x \forall y \ (ISA(x,y) \Leftrightarrow \exists t \ \Pi_P(x,t,y))$$

This mirrors the fact that, in any MINE Γ , $x \to y$ may be inferred if and only if a positive path is derivable that has x as initial node and y as terminal node. The predicate Π_N is defined in a completely analogous way to Π_P .

This concludes the definition of LT, a FOL theory that may then be concisely described by means of the following definition.

Definition 8 The Logical Theory of Multiple Inheritance with Exceptions (LT) is the FOL theory $LT \equiv \{A_1, A_2, A_3, A_4\}$, where:

•
$$A_1 \equiv \forall x \forall t \forall y \ \Pi_P(x, t, y) \Leftrightarrow$$

$$((x = t \land ISDA(x, y) \land \neg ISDNA(x, y)) \land$$

$$(((\exists t'\Pi_P(x, t, t') \land t \neq t' \land ISDA(t', y)) \lor$$

$$(\exists t'\Pi_P(x, t', t) \land t \neq t' \land ISDA(t, y))) \land$$

$$\neg ISDNA(x, y) \land$$

$$(\forall v \exists v'(\Pi_P(x, v', v) \land ISDNA(v, y))$$

$$\Rightarrow (\exists z ISDA(z, y) \land (z = x \lor \Pi_P(x, z, v))))))$$

$$A_{2} \equiv \forall x \forall t \forall y \ \Pi_{N}(x,t,y) \Leftrightarrow \\ ((x = t \land ISDNA(x,y) \land \neg ISDA(x,y)) \\ \lor \ (((\exists t'\Pi_{P}(x,t,t') \land t \neq t' \land ISDNA(t',y)) \lor \\ (\exists t'\Pi_{P}(x,t',t) \land t \neq t' \land ISDNA(t,y))) \land \\ \neg ISDA(x,y) \land \\ (\forall v \exists v'(\Pi_{P}(x,v',v) \land ISDA(v,y)) \\ \Rightarrow (\exists z ISDNA(z,y) \land (z = x \lor \Pi_{P}(x,z,v)))))))$$

- $A_3 \equiv \forall x \forall y \ (ISA(x,y) \Leftrightarrow \exists t \ \Pi_P(x,t,y));$
- $A_4 \equiv \forall x \forall y \ (ISNA(x,y) \Leftrightarrow \exists t \ \Pi_N(x,t,y))$

4.3 An equivalence result

So far we have described $\mathcal{M}(\Gamma)$ and LT, the two components of a (non-monotonic) function that maps MINEs into sets of FOL formulae. At this point the only thing we have to do is to describe the properties of $\mathcal{T}(\Gamma) \equiv \mathcal{M}(\Gamma) \cup LT$. By exploiting the acyclicity of our MINEs, we are able to prove the following fundamental lemma by finite induction on the degree of paths.

Lemma 1 (Correspondence) Let Γ be a MINE and \mathbf{M} be any first order interpretation such that $\mathbf{M} \models_{LT} \mathcal{M}(\Gamma)$. Then, for all $\langle x, y \rangle \subseteq (\Gamma_I \cup \Gamma_C) \times (\Gamma_I \cup \Gamma_C)$:

1.
$$x \to y \in C(\Gamma) \iff \mathbf{M} \models_{LT} ISA(x, y)$$

2. $x \neq y \in C(\Gamma) \iff \mathbf{M} \models_{LT} ISNA(x, y)$

The proofs of Lemma 1 and of the lemmas that follow are given in the appendix. The following lemma derives immediately from Lemma 1.

Lemma 2 (Completion) Let Γ be a MINE and \mathbf{M} be any first order interpretation such that $\mathbf{M} \models_{LT} \mathcal{M}(\Gamma)$. Then, for all $\langle x, y \rangle \subseteq (\Gamma_I \cup \Gamma_C) \times (\Gamma_I \cup \Gamma_C)$:

1.
$$x \to y \notin C(\Gamma) \iff \mathbf{M} \models_{LT} \neg ISA(x, y)$$

2. $x \neq y \notin C(\Gamma) \iff \mathbf{M} \models_{LT} \neg ISNA(x, y)$

By combining Lemma 1, Lemma 2 and the Consistency Lemma for " \vdash " proven in [6] we obtain

Lemma 3 (Consistency) Let Γ be a MINE and let \mathbf{M} be any first order interpretation such that $\mathbf{M} \models_{LT} \mathcal{M}(\Gamma)$. Then, for all $\langle x, y \rangle \subseteq (\Gamma_I \cup \Gamma_C) \times (\Gamma_I \cup \Gamma_C)$:

1.
$$\mathbf{M} \models_{LT} ISA(x, y) \Longrightarrow \mathbf{M} \models_{LT} \neg ISNA(x, y)$$

2.
$$\mathbf{M} \models_{LT} ISNA(x, y) \Longrightarrow \mathbf{M} \models_{LT} \neg ISA(x, y)$$

Hence, this lemma shows the correct relation that exists between ISA and ISNA.

The results dealt with in the preceding lemmas presuppose the existence of at least one model of $\mathcal{M}(\Gamma)$ with respect to LT. We can prove that, in fact, such a model always exists. The following lemma is proven in a constructive way, i.e. by deriving the required model from Γ itself.

Lemma 4 Let Γ be a MINE. Then there exists a first order interpretation \mathbf{M} such that $\mathbf{M} \models_{LT} \mathcal{M}(\Gamma)$.

Finally, we may summarize the results of Lemmas 1-4 into the following theorem that completely characterizes $\mathcal{M}(\Gamma) \models_{LT}$.

Theorem 1 Let Γ be a MINE. Then the following propositions hold

1.
$$x \to y \in C(\Gamma) \iff \mathcal{M}(\Gamma) \models_{LT} ISA(x, y)$$

2. $x \not\to y \in C(\Gamma) \iff \mathcal{M}(\Gamma) \models_{LT} ISNA(x, y)$
3. $x \to y \notin C(\Gamma) \iff \mathcal{M}(\Gamma) \models_{LT} \neg ISA(x, y)$
4. $x \not\to y \notin C(\Gamma) \iff \mathcal{M}(\Gamma) \models_{LT} \neg ISNA(x, y)$
5. $\mathcal{M}(\Gamma) \models_{LT} ISA(x, y) \Longrightarrow \mathcal{M}(\Gamma) \models_{LT} \neg ISNA(x, y)$

6.
$$\mathcal{M}(\Gamma) \models_{LT} ISNA(x, y) \Longrightarrow \mathcal{M}(\Gamma) \models_{LT} \neg ISA(x, y)$$

Clauses (1) and (2) of Theorem 1 show that \mathcal{T} is a *complete* translation of MINEs into FOL, in the sense that, for each conclusion that can be derived from a MINE Γ , an equivalent conclusion is derivable from $\mathcal{T}(\Gamma)$.

They also show that the translation is *sound*, in the sense that, for each conclusion that may be derived from $\mathcal{T}(\Gamma)$, an equivalent conclusion is derivable from Γ itself. However, this latter statement must be interpreted with a *caveat*: some formulae that can be derived from $\mathcal{T}(\Gamma)$ have in fact no equivalent conclusions derivable from Γ , in the sense that their would-be equivalents are not even expressible in the MINE formalism. For instance, in the case of a MINE Γ such that $C(\Gamma)$ comprises the edge $p \to q$, the set of formulae derivable from $\mathcal{T}(\Gamma)$ will indeed contain the formula ISA(p,q), but will also contain formulae that do not correspond to conclusions of Γ : among these, formulae containing connectives (such as e.g. $ISA(p,q) \wedge ISA(p,q)$), formulae containing instances of predicates different from ISAand ISNA (such as e.g. $\exists t \prod_{P}(p,t,q)$), tautologies of the first order predicate calculus, etc. This is an obvious consequence of the fact that we are dealing with a translation of a formalism into an expressively richer one, and does not invalidate the substance of our claim; because of this we have substantially ignored the issue elsewhere in this paper.

A related comment may be made for clauses (3) and (4) of Theorem 1, in that formulae of the form $\neg ISA(x, y)$ and $\neg ISNA(x, y)$ do not find a direct analogue in the MINE formalism: the result is that the non-derivability of a conclusion (e.g. it is not derivable that birds are typically yellow) finds its translation in the derivability of the negated conclusion (e.g. it is not the case that birds are typically yellow). This can be seen as a *completion* of the theory with respect to the *ISA* and *ISNA* predicates.

In order to better understand the properties of the \mathcal{T} translation, we end this section by reporting a few examples.

Example 1 Let Γ be a MINE such that $\Gamma_A = \{o \to P, P \to B, B \to F\}$, a MINE that encodes our previous "birds" example. The edge $o \to F$ belongs to $C(\Gamma)$; in keeping with this, the formula ISA(o, F) is a logical consequence of $\mathcal{T}(\Gamma)$. If we add the edge $P \not\to F$ to Γ_A , the edge $o \to F$ no more belongs to $C(\Gamma)$, which instead includes $o \not\to F$; accordingly, the formula ISA(o, F) is no more a logical consequence of $\mathcal{T}(\Gamma)$, which instead now includes ISNA(o, F). This shows that the \mathcal{T} function is non-monotonic.

Example 2 Let Γ be a MINE such that $\Gamma_A = \{n \to Q, n \to R, Q \to P, R \neq P\}$, a MINE that encodes the famous "Nixon diamond". Neither $n \to P$ nor $n \neq P$ belong to $C(\Gamma)$, corresponding to the fact that, according to the "skeptical" approach, whether Nixon is a pacifist or not cannot reasonably be concluded from the information contained in an inherently ambiguous network such as Γ . Accordingly, the formula $\neg ISA(n, P) \land \neg ISNA(n, P)$ is a logical consequence of $\mathcal{T}(\Gamma)$ (this is a direct consequence of Lemma 2).

Example 3 Let Γ be a MINE such that $\Gamma_A = \{r \to N, r \not\to R, N \to Q, N \to R, Q \to P, R \not\to P\}$, a MINE that encodes an extended version of the "Nixon diamond" (where the

predicate N may be interpreted as "member of the Nixon family"), reported as Figure 5 of [19]. Both $r \to Q$ and $r \to P$ belong to $C(\Gamma)$, corresponding to the fact that the ambiguity whether Richard is a pacifist or not can now be resolved thanks to the information that Richard is not a republican, although typical Nixons are. Accordingly, both ISA(r, Q) and ISA(r, P) are logical consequences of $\mathcal{T}(\Gamma)$.

5 Conclusion

We have described an approach to non-monotonic reasoning that combines the advantages (in terms of incrementality of the knowledge acquisition phase) that are offered by MINE representation languages, with those (in terms of semantic clarity) that are offered by a formally solid apparatus such as First Order Logic. We have achieved this by specifying a first order representation LT of the MINE formalism, and a mapping function \mathcal{M} that maps a MINE Γ into a set of FOL formulae. In our framework, a KB is obtained by adding to LT the result of the application of \mathcal{M} to a particular MINE, thus obtaining the FOL knowledge base $\mathcal{T}(\Gamma) = \mathcal{M}(\Gamma) \cup LT$. The net effect is that the LT component of $\mathcal{T}(\Gamma)$ remains fixed throughout the phase of knowledge acquisition, while the only component that varies is the $\mathcal{M}(\Gamma)$ component. We have shown how this variation is achieved by performing simple piecemeal additions of new chunks of knowledge to the pre-existing KB, thereby accomplishing incrementality and compositionality. We have also shown that \mathcal{T} is a sound and complete mapping, as each conclusion derivable from a MINE Γ is also derivable from $\mathcal{T}(\Gamma)$, and vice-versa.

An interesting side-effect of this result is that MINE-like reasoning can be performed by means of the ordinary and widely available first order theorem provers. A simple interface may be devised that allows the user to convey his information to the KB in terms of the simpler MINE representation language; it is then the interface that takes the charge of implementing the \mathcal{T} translation function, performing the required additive modifications to axioms (8),(9) and (10). We are currently exploring the possibility of an efficient implementation of this framework by using the "theory resolution" technique for first order theorem proving devised by Stickel [17], which allows us to "wire" the *LT* theory into a resolution-based theorem prover, thereby freeing the knowledge base management system from the need to handle the axioms of *LT* directly.

References

- [1] Proceedings of the Third International Workshop on Nonmonotonic Reasoning, Lake Tahoe, CA, 1990.
- [2] Douglas E. Appelt and Kurt Konolige. A nonmonotonic logic for reasoning about speech acts and belief revision. In Michael Reinfrank, Johan de Kleer, Matthew L. Gin-

sberg, and Erik Sandewall, editors, *Nonmonotonic reasoning*, pages 164–175. Springer, Heidelberg, FRG, 1989.

- [3] Nino Cocchiarella. Predication versus membership in the distinction between logic as language and logic as calculus. *Synthese*, 77:37–72, 1988.
- [4] Matthew L. Ginsberg, editor. Readings in nonmonotonic reasoning. Morgan Kaufmann, Los Altos, CA, 1987.
- [5] John F. Horty and Richmond H. Thomason. Mixing strict and defeasible inheritance. In Proceedings of AAAI-88, 7th Conference of the American Association for Artificial Intelligence, pages 427–432, St. Paul, MN, 1988.
- [6] John F. Horty, Richmond H. Thomason, and David S. Touretzky. A skeptical theory of inheritance in nonmonotonic semantic networks. In *Proceedings of AAAI-87, 6th Conference of the American Association for Artificial Intelligence*, pages 358–363, Seattle, WA, 1987. [a] An extended version appears as [7].
- [7] John F. Horty, Richmond H. Thomason, and David S. Touretzky. A skeptical theory of inheritance in nonmonotonic semantic networks. *Artificial Intelligence*, 42:311–348, 1990.
- [8] Maurizio Lenzerini, Daniele Nardi, and Maria Simi, editors. *Inheritance hierarchies in knowledge representation and programming languages*. Wiley, Chichester, UK, 1991.
- [9] Vladimir Lifschitz. Computing circumscription. In Proceedings of IJCAI-85, 9th International Joint Conference on Artificial Intelligence, pages 121–127, Los Angeles, CA, 1985. [a] Also reprinted in [4], pp. 167–173.
- [10] Vladimir Lifschitz. Towards a metatheory of action. In Proceedings of KR-91, 2nd International Conference on Principles of Knowledge Representation and Reasoning, pages 376–386, Cambridge, MA, 1991.
- [11] Alan K. Mackworth and Raymond Reiter. A logical framework for depiction and image interpretation. Artificial Intelligence, 41:125–155, 1989.
- [12] Michael Reinfrank, Johan de Kleer, Matthew L. Ginsberg, and Erik Sandewall, editors. Nonmonotonic reasoning. Springer, Heidelberg, FRG, 1989.
- [13] Raymond Reiter. A theory of diagnosis from first principles. Artificial Intelligence, 32:57–95, 1987. [a] Also reprinted in [4], pp. 352–371.
- [14] Erik Sandewall. Nonmonotonic inference rules for multiple inheritance with exceptions. Proceedings of the IEEE, 13:1345–1353, 1986.
- [15] Fabrizio Sebastiani. On heterogeneous model-preference default theories. In Proceedings of CSCSI/SCEIO-90, 8th Biennial Conference of the Canadian Society for Computational Studies of Intelligence, pages 84–91, Ottawa, Ontario, 1990.

- [16] Bart Selman and Hector J. Levesque. The tractability of path-based inheritance. In Proceedings of IJCAI-89, 11th International Joint Conference on Artificial Intelligence, pages 1140–1145, Detroit, MI, 1989. [a] Also reprinted in [8], pp. 83–96.
- [17] Mark E. Stickel. Automated deduction by theory resolution. Journal of Automated Reasoning, 1:333–355, 1985.
- [18] David S. Touretzky. The mathematics of inheritance systems. Pitman, London, UK, 1986.
- [19] 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.

A Proofs

Proof of Lemma 1

Hereafter, individual constants will sometimes be called "nodes", given the biunivocal correspondence between $\Gamma_I \cup \Gamma_C$ and the set of individual constants of our FOL language. We will allow the use of the symbol $deg_{\Gamma}(x, y)$ instead of $deg_{\Gamma}(\sigma)$ whenever σ is a path with initial node x and terminal node y.

We will only prove clause (1), as the proof of clause (2) is completely analogous. Let us consider the right-hand side of axiom A_1 in Definition 8, deprived of its top-level connectives:

$$x = t \land ISDA(x, y) \land \neg ISDNA(x, y)$$

$$((\exists t' \Pi_{P}(x, t, t') \land t \neq t' \land ISDA(t', y))$$

$$(11)$$

$$(\exists t' \Pi_P(x, t', t) \land t \neq t' \land ISDA(t, y)))$$

$$\forall v \exists v' (\Pi_P(x, t', t) \land t \neq t' \land ISDA(t, y)))$$

$$(12)$$

$$\Rightarrow (\exists z \ ISDA(z, y) \land (z = x \lor \Pi_P(x, z, v))$$
(13)

$$\neg ISDNA(x,y)$$
 (14)

That is, the right-hand side of axiom A_1 in Definition 8 is exactly $(11) \lor ((12) \land (13) \land (14))$.

(\Leftarrow) We show that, if $\mathbf{M} \models ISA(x, y)$, then $x \to y \in C(\Gamma)$. As $\mathbf{M} \models \mathcal{T}(\Gamma)$ and $\mathbf{M} \models ISA(x, y)$, then, by axiom A_3 in Definition 8, there exists a node t_1 such that $\mathbf{M} \models \Pi_P(x, t_1, y)$. We show by induction on $n = deg_{\Gamma}(x, y)$ that, if $\mathbf{M} \models \Pi_P(x, t_1, y)$, then $\Gamma \vdash x \to \delta_1 \to t_1 \to \delta_2 \to y$ for some δ_1 and δ_2 , hence $x \to y \in C(\Gamma)$.

Case n=1: Let $\mathbf{M} \models \Pi_P(x, t_1, y)$, and let $deg_{\Gamma}(x, y) = 1$. By axiom A_1 in Definition 8, \mathbf{M} is a model of (11), while it is not a model of the conjunction of

(12), (13) and (14). Since **M** is a model of (11), it follows that $x = t_1$ and $\mathbf{M} \models ISDA(x, y) \land \neg ISDNA(x, y)$; hence, because of the biunivocal correspondence between the edges of Γ and the instances of ISDA and ISDNA, it also follows that $x \to y \in \Gamma_A$ and $x \neq y \notin \Gamma_A$. From Definition 4, it follows that $\Gamma \vdash x \to y$.

It is worth to observe that, if we had not minimized the predicates ISDA and ISDNA by means of axioms (8) and (9), there would not be a biunivocal correspondence between the edges of Γ and the instances of ISDA and ISDNA. Without the minimization of predicates ISDA and ISDNA that is enforced by axioms (8) and (9), if e.g. $\mathbf{M} \models ISDA(x, y)$, then it would not necessarily follow that $x \to y \in \Gamma_A$; hence, in our case, from $\mathbf{M} \models ISDA(x, y) \land \neg ISDNA(x, y)$ it would not necessarily follow that $x \to y \in \Gamma_A$ and $x \not\to y \notin \Gamma_A$; as a consequence, Lemma 1 would not hold.

inductive step: Let the thesis hold for $deg_{\Gamma}(x, y) < n$; let us show that it also holds for $deg_{\Gamma}(x, y) = n$. Let $\mathbf{M} \models \Pi_{P}(x, t_{1}, y)$, and let $deg_{\Gamma}(x, y) = n$. By axiom A_{1} in Definition 8, if \mathbf{M} is a model of (11), then, similarly to the case of n = 1, it is true that $\Gamma \vdash x \rightarrow y$. If, instead, \mathbf{M} is not a model of (11), then \mathbf{M} must be a model of the conjunction of (12), (13) and (14). Hence, with no loss of generality, from (12) we may say that there exists a node $t_{2} \neq t_{1}$ such that $\mathbf{M} \models \Pi_{P}(x, t_{2}, t_{1}) \land ISDA(t_{1}, y)$; from (14), it is true that $\mathbf{M} \models \neg ISDNA(x, y)$; last, from (13), it is true that for any node v_{1} for which there exists a node v_{2} such that $\mathbf{M} \models \Pi_{P}(x, v_{2}, v_{1}) \land ISDNA(v_{1}, y)$, there exists a node z_{1} such that $\mathbf{M} \models ISDA(z_{1}, y)$ and either $z_{1} = x$ or $\mathbf{M} \models \Pi_{P}(x, z_{1}, v_{1})$. Because of the biunivocal correspondence between the edges of Γ and the instances of ISDA and ISDNA, there are edges $t_{1} \rightarrow y \in \Gamma_{A}$, $x \neq y \notin \Gamma_{A}$, $v_{1} \not\rightarrow y \in \Gamma_{A}$ and $z_{1} \rightarrow y \in \Gamma_{A}$. Hence, it is true that $deg_{\Gamma}(x, t_{1}) < deg_{\Gamma}(x, y)$, as it would otherwise be true that $deg_{\Gamma}(x, t_{1}) \ge deg_{\Gamma}(x, y) \ge deg_{\Gamma}(x, t_{1}) + 1$. Analogously, it is true that $deg_{\Gamma}(x, v_{1}) < deg_{\Gamma}(x, y)$. Hence, for inductive hypothesis on $\mathbf{M} \models \Pi_{P}(x, t_{2}, t_{1})$, $\mathbf{M} \models \Pi_{P}(x, v_{2}, v_{1})$ and $\mathbf{M} \models \Pi_{P}(x, z_{1}, v_{1})$, the following holds:

- 1. $t_1 \rightarrow y \in \Gamma_A$
- 2. $\Gamma \vdash x \to \sigma_1 \to t_1$
- 3. $x \not\rightarrow y \notin \Gamma_A$
- 4. for all v_1 and for all σ_2 such that $\Gamma \vdash x \to \sigma_2 \to v_1$ and $v_1 \not\to y \in \Gamma_A$, there exists a z_1 such that $z_1 \to y \in \Gamma_A$ and either $z_1 = x$ or $\Gamma \vdash x \to \gamma_1 \to z_1 \to \gamma_2 \to v_1$ for some path γ_1 and for some path γ_2 .

Hence, from Definition 4, it is true that $\Gamma \vdash x \to \sigma_1 \to t_1 \to y$.

(⇒) Let $x \to y \in C(\Gamma)$. We show that $\mathbf{M} \models ISA(x, y)$. From $x \to y \in C(\Gamma)$ it follows that $\Gamma \vdash x_1 \to x_2 \to \ldots \to x_{n-1} \to x_n$, for some nodes x_1, \ldots, x_n , with $x_1 = x, x_n = y$ and n > 1. Let us show, by induction on $m = deg_{\Gamma}(x_1 \to \ldots \to x_n)$, that $\mathbf{M} \models \prod_P(x_1, x_j, x_n), 1 \leq j < n$, hence $\mathbf{M} \models ISA(x, y)$. **Case m=1:** From $\Gamma \vdash x_1 \to x_2$ we obtain $x_1 \to x_2 \in \Gamma_A$ and $x_1 \not\to x_2 \notin \Gamma_A$. Hence, $\mathbf{M} \models ISDA(x_1, x_2) \land \neg ISDNA(x_1, x_2)$, and, from (11), $\mathbf{M} \models \Pi_P(x_1, x_1, x_2)$.

inductive step: Let the thesis be true for any path σ' such that $deg_{\Gamma}(\sigma') < m$; let us show that it is also true for paths $\sigma = x_1 \to \ldots \to x_n$ for which $deg_{\Gamma}(\sigma') = m$, with m > 1. If n = 2, then, similarly to the case m = 1, we obtain that $\mathbf{M} \models \Pi_P(x_1, x_1, x_2)$. Otherwise, from $\Gamma \vdash x_1 \to x_2 \to \ldots \to x_{n-1} \to x_n$, with n > 2 we obtain:

- 1. $x_{n-1} \to x_n \in \Gamma_A$
- 2. $\Gamma \vdash x_1 \rightarrow \ldots \rightarrow x_{n-1}$
- 3. $x_1 \not\rightarrow x_n \notin \Gamma_A$
- 4. for all v and for all σ_1 such that $\Gamma \vdash x_1 \to \sigma_1 \to v$ and $v \not\to x_n \in \Gamma_A$ there exists z such that $z \to x_n \in \Gamma_A$ and either $z = x_1$ or $\Gamma \vdash x_1 \to \gamma_1 \to z_1 \to \gamma_2 \to v$ for some path γ_1 and for some path γ_2

For inductive hypothesis on $x_1 \to \ldots \to x_{n-1}$, $\mathbf{M} \models \Pi_P(x_1, x_j, x_{n-1})$ holds for any $1 \leq j < n-1$. Besides, $\mathbf{M} \models ISDA(x_{n-1}, x_n)$ and $\mathbf{M} \models \neg ISDNA(x_1, x_n)$. Hence, \mathbf{M} is a model for (13). In fact, for any node v_1 for which there exists a node v_2 such that $\mathbf{M} \models \Pi_P(x_1, v_2, v_1)$ and $\mathbf{M} \models ISDNA(v_1, x_n)$, from the (\Leftarrow) we obtain that $\Gamma \vdash x_1 \to \sigma_1 \to v_1$ and $v_1 \neq x_n \in \Gamma_A$; hence there exists a z such that $z \to x_n \in \Gamma_A$ and either $z = x_1$ or $\Gamma \vdash x_1 \to \gamma_1 \to z \to \gamma_2 \to v_1$ for some γ_1 and γ_2 . But in this case, for inductive hypothesis, we obtain $\mathbf{M} \models ISDA(z, x_n) \land (\Pi_P(x_1, z, v_1) \lor z = x_1)$. It follows that \mathbf{M} is a model of (12), (13) and (14), hence $\mathbf{M} \models \Pi_P(x_1, x_j, x_n), 1 \leq j < n$.

Proof of Lemma 2

Straightforward.

Proof of Lemma 3

Only clause (1) needs to be proven, as the proof for clause (2) is completely analogous. From $\mathbf{M} \models ISA(x, y)$ we obtain, by Lemma 1, that $x \to y \in C(\Gamma)$. Hence, from the Consistency Theorem for " \vdash " presented in [6], we obtain that $x \neq y \notin C(\Gamma)$; hence, from Lemma 2, $\mathbf{M} \models \neg ISNA(x, y)$ holds.

Proof of Lemma 4

Let us build a model \mathbf{M} of $\mathcal{M}(\Gamma)$ with respect to LT in the following way. Let the domain D be given by $D = \Gamma_I \cup \Gamma_C$; besides, for any $x \in \Gamma_I \cup \Gamma_C$, let $\mathbf{M}(x) = x$. Let the following conditions hold for any set $\{x, t, y\} \subseteq D$:

- $\mathbf{M}(ISDA)(x, y) = true$ if and only if $x \to y \in \Gamma_A$
- $\mathbf{M}(ISDNA)(x, y) = true$ if and only if $x \not\rightarrow y \in \Gamma_A$
- $\mathbf{M}(ISA)(x, y) = true$ if and only if $x \to y \in C(\Gamma)$
- $\mathbf{M}(ISNA)(x, y) = true$ if and only if $x \not\rightarrow y \in C(\Gamma)$
- $\mathbf{M}(\Pi_P)(x, x, y) = true$ if and only if $\Gamma \vdash x \to y$
- $\mathbf{M}(\Pi_N)(x, x, y) = true$ if and only if $\Gamma \vdash x \not\rightarrow y$
- $\mathbf{M}(\Pi_P)(x, t, y) = true$ if and only if $\Gamma \vdash x \to \tau_1 \to t \to \tau_2 \to y$
- $\mathbf{M}(\Pi_N)(x, t, y) = true$ if and only if $\Gamma \vdash x \to \tau_1 \to t \to \tau_2 \not\to y$

It is straightforward to show that this interpretation is a model of $\mathcal{M}(\Gamma)$ with respect to LT.