Typicality-based revision for handling exceptions in Description Logics

Roberto Micalizio[1⋆] and Gian Luca Pozzato[1⋆⋆]

Dip. Informatica - Universitá di Torino - Italy
{roberto.micalizio,gianluca.pozzato}@unito.it

Abstract. We continue our investigation on how to revise a Description Logic knowledge base when detecting exceptions. Our approach relies on the methodology for debugging a Description Logic terminology, addressing the problem of diagnosing inconsistent ontologies by identifying a minimal subset of axioms responsible for an inconsistency. In the approach we propose, once the source of the inconsistency has been localized, the identified TBox inclusions are revised in order to obtain a consistent knowledge base including the detected exception. We define a revision operator whose aim is to replace inclusions of the form "Cs are Ds" with "typical Cs are Ds", admitting the existence of exceptions, obtaining a knowledge base in the nonmonotonic logic $\mathcal{ALC}^{\mathbf{R}}_{min}\mathbf{T}$ which corresponds to a notion of rational closure for Description Logics of typicality. We also describe an algorithm implementing such a revision operator.

1 Introduction

The family of Description Logics (DL) [1] is one of the most important formalisms of knowledge representation. DLs are reminiscent of the early semantic networks and of frame-based systems. They offer two key advantages: (i) a well-defined semantics based on first-order logic, and (ii) a good trade-off between expressivity and computational complexity. DLs have been successfully implemented by a range of systems, and are at the base of languages for the Semantic Web such as OWL. In a DL framework, a knowledge base (KB) comprises two components: (i) a TBox, containing the definition of concepts (and possibly roles) and a specification of inclusion relations among them; and (ii) an ABox, containing instances of concepts and roles, in other words, properties and relations between individuals.

Recent works [2–7] have addressed the problem of diagnosing inconsistent ontologies by identifying a minimal subset of axioms responsible for the inconsistency. The idea of these works is that once the source of the inconsistency has been localized, the ontology engineer can intervene and revise the identified axioms by rewriting or removing some of them in order to restore the consistency. These approaches presuppose that the ontology has become inconsistent due to the introduction of *errors*; as for instance when two ontologies are merged together.

⋆ The author is partially supported by the project "AThOS: Accountable Trustworthy Organizations and Systems" , Università di Torino and Compagnia di San Paolo.

⋆⋆ The author is partially supported by the project "ExceptionOWL: Nonmonotonic Extensions of Description Logics and OWL for defeasible inheritance with exceptions" , Università di Torino and Compagnia di San Paolo.

In this work we continue our investigation in revising a Description Logic knowledge base when an *exception* is discovered. Albeit an exception has the same effect of an error (i.e., it causes an ontology to become inconsistent), an exception is not an error. Rather, an exception is a piece of additional knowledge that, although partially in contrast with what we know, must be taken into account. Thus, on the one hand, ignoring exceptions would be deleterious as the resulting ontology would not reflect the applicative domain correctly. On the other hand, accommodating exceptions requires the exploitation of some form of defeasible reasoning that allows us to revise some of concepts in the ontology.

In [8] we have moved a first step in the direction of providing a methodology for revising a TBox when a new concept is introduced and the resulting TBox is *incoherent*, i.e. it contains at least a concept whose interpretation is mapped into an empty set of domain elements. Here we refine that proposal, and move a further step in order to tackle the problem of revising a TBox in order to accommodate a newly received information about an exception represented by an ABox individual x. Our approach is inspired by the weakening-based revision introduced in [9] and relies on the methodology by Schlobach et al. [2–4] for detecting exceptions by identifying a minimal subset of axioms responsible for an inconsistency. Once the source of the inconsistency has been localized, the identified axioms are revised in order to obtain a consistent knowledge base including the detected exception about the individual x. To this aim, we use a nonmonotonic extension of the DL $\mathcal{ALC}$ recently introduced by Giordano and colleagues in [10]. This extension is based on the introduction of a typicality operator $\mathbf{T}$ in order to express typical inclusions. The intuitive idea is to allow concepts of the form $\mathbf{T}(C)$, whose intuitive meaning is that $\mathbf{T}(C)$ selects the *typical* instances of a concept C. It is therefore possible to distinguish between properties holding for all instances of a concept C ($C \sqsubseteq D$), and those holding only for the typical instances of C ($\mathbf{T}(C) \sqsubseteq D$). For instance, a knowledge base can consistently express that birds normally fly ($\mathbf{T}(Bird) \sqsubseteq Fly$), but penguins are exceptional birds that do not fly ($Penguin \sqsubseteq Bird$ and $Penguin \sqsubseteq \neg Fly$). The $\mathbf{T}$ operator is intended to enjoy the well-established properties of *rational logic*, introduced by Lehmann and Magidor in [11] for propositional logic. In order to reason about prototypical properties and defeasible inheritance, the semantics of this nonmonotonic DL, called $\mathcal{ALC}^{\mathbf{R}}_{min}\mathbf{T}$, is based on rational models and exploits a minimal models mechanism based on the minimization of the rank of the domain elements. This semantics corresponds to a natural extension to DLs of Lehmann and Magidor's notion of *rational closure* [11].

Given a consistent knowledge base $K = (\mathcal{T}, \mathcal{A})$ and a consistent ABox $\mathcal{A}' = \{D_1(x), D_2(x), \ldots, D_n(x)\}$, such that $(\mathcal{T}, \mathcal{A} \cup \mathcal{A}')$ is inconsistent, we define a typicality-based revision of $\mathcal{T}$ in order to replace some inclusions $C \sqsubseteq D$ in $\mathcal{T}$ with $\mathbf{T}(C) \sqsubseteq D$, resulting in a new TBox $\mathcal{T}^{new}$ such that $(\mathcal{T}^{new}, \mathcal{A} \cup \mathcal{A}')$ is consistent in the nonmonotonic $\mathcal{ALC}^{\mathbf{R}}_{min}\mathbf{T}$ and such that $\mathcal{T}^{new}$ captures a notion of *minimal changes*. As an example, consider a knowledge base $K = (\mathcal{T}, \mathcal{A})$ whose TBox $\mathcal{T}$ is as follows:

$$Professor \sqsubseteq AcademicStaffMember$$
$$Professor \sqsubseteq \exists teaches.Course$$

representing that professors are members of the academic staff that teach courses, and whose ABox contains the information that Frank is a professor, namely $\mathcal{A} =$

$\{Professor(frank)\}$. If further information $\mathcal{A}'$ about Frank is provided, for instance that he does not teach any course because he asked for a sabbatical year:

$$\mathcal{A}' = \{\neg\exists teaches.Course(frank)\}$$

the resulting knowledge base $(\mathcal{T}, \mathcal{A} \cup \mathcal{A}')$ is inconsistent. Our approach, given the exception arised in $\mathcal{A}'$, provides a typicality-based revision of $\mathcal{T}$ as described by the following $\mathcal{T}^{new}$:

$Professor \sqsubseteq AcademicStaffMember$
$\mathbf{T}(Professor) \sqsubseteq \exists teaches.Course$

representing that, in normal circumstances, professors teach courses, admitting the existence of exceptions. The revised TBox is now consistent with all the information of the ABox, namely the knowledge base $(\mathcal{T}^{new}, \mathcal{A} \cup \mathcal{A}')$ is consistent in the DL of typicality $\mathcal{ALC}^{\mathbf{R}}_{min}\mathbf{T}$.

2 State of the Art and Motivations

Our work starts from the contribution in [9], where a weakening-based revision operator is introduced in order to handle exceptions in an $\mathcal{ALC}$ knowledge base. In that work, the authors consider the following problem: given a consistent $\mathcal{ALC}$ knowledge base K, further information is provided by an additional, consistent $\mathcal{ALC}$ knowledge base K', however the resulting $K \cup K'$ is inconsistent. K' contains only ABox assertions that, as a matter of fact, correspond to exceptions. The basic idea of [9] is that conflicts are solved by adding explicit exceptions to weaken them, and then assuming that the number of exceptions is minimal. A revision operator $\circ_w$ is introduced in order to weaken the initial knowledge base K to handle the exceptions in K': in other words, a *disjunctive* knowledge base $K \circ_w K' = K_1, K_2, \ldots, K_n$, where each K_i is such that $K \cup K_i$ is consistent, represents the revised knowledge. Each K_i is obtained from K as follows: if K' contains information about an exception x; e.g., $\neg D(x) \in K'$ when $C(x)$ and $C \sqsubseteq D \in K$, then either the inclusion relation is replaced by $C \sqcap \neg\{x\} \sqsubseteq D$ or $C(x)$ is replaced by $\top(x)$. Then, the revision operator $\circ_w$ selects the TBoxes minimizing a degree of the weakened knowledge base, defined as the number of exceptions introduced. A further refined revision operator, taking into account exceptions introduced by universal restrictions, is introduced.

Let us consider the following example. Let K contain $Bird \sqsubseteq Fly$ and $Bird(tweety)$. Let K' be the newly introduced ABox, containing the only information about the fact that Tweety does not fly (because he is, for instance, a penguin), so $K' = \{\neg Fly(tweety)\}$. The revision operator $\circ_w$ introduces the following two weakening-based knowledge bases:

$K_1 = \{Bird \sqcap \neg\{tweety\} \sqsubseteq Fly\}$
$K_2 = \{\top(tweety)\}$

Let us further consider the well known example of the Nixon diamond: let $K = (\mathcal{T}, \emptyset)$ where $\mathcal{T}$ is:

$Quacker \sqsubseteq \mathit{Pacifist}$
$Republican \sqsubseteq \neg \mathit{Pacifist}$

and let $K' = \{Quacker(nixon), Republican(nixon)\}$. In this case, we have that $K \circ_w K' = \{K_1, K_2\}$, where $K_1 = \{Quacker \sqcap \neg\{nixon\} \sqsubseteq \mathit{Pacifist}, Republican \sqsubseteq \neg \mathit{Pacifist}\}$ and $K_2 = \{Quacker \sqsubseteq \mathit{Pacifist}, Republican \sqcap \neg\{nixon\} \sqsubseteq \neg \mathit{Pacifist}\}$. On the contrary, the knowledge base obtained by weakening both the inclusions is discarded, since the degree of exceptionality is higher with respect to the alternatives: in both K_1 and K_2 only one exception is introduced, whereas the discarded alternative weaken two inclusions.

The above simple examples show the drawbacks of the revision operator introduced in [9]. In the first example, it can be observed that no inferences about Tweety can be further performed in the revised knowledge base when $Bird(tweety)$ is weakened to $\top(tweety)$, because Tweety is no longer considered as a bird: for instance, if K further contains the inclusion $Bird \sqsubseteq Animal$, we are no longer able to infer that Tweety is an animal, i.e. $Animal(tweety)$. Furthermore, in the weakened inclusion $Bird \sqcap \neg\{tweety\} \sqsubseteq Fly$ all exceptions are explicitly enumerated, however this implies that a revision would be needed if further exceptions are discovered, in other words if we further know that also *cipcip* is a not flying bird, the resulting knowledge base is still inconsistent, and we need to revise it by replacing the inclusion relation with $Bird \sqcap \neg\{tweety\} \sqcap \neg\{cipcip\} \sqsubseteq Fly$. A similar problem also occurs in the example of the Nixon diamond: since a disjunctive knowledge base $\{K_1, K_2, \dots, K_n\}$ is satisfied in a model $\mathcal{M}$ if and only if $\mathcal{M}$ is a model of at least one K_i, all K_is become inconsistent when another individual being both a Quacker and a Republican is discovered.

The approach in [9] is strongly related to well established works in inconsistency handling in knowledge bases formalized in propositional and first-order logics. In particular, basic ideas are inspired to the DMA (disjunctive maxi-adjustment) approach proposed in [12] in order to handle conflicts in stratified propositional knowledge bases. This approach is extended to first order knowledge bases in [13], whose basic idea is similar to the one of weakening an inclusion: here a first-order formula $\forall x P(x) \rightarrow Q(x)$ generating a conflict is weakened by dropping the instances originating such inconsistence, rather than dropping the whole formula. An alternative approach, called RCMA (Refined Conjunctive Maxi-Adjustment), is proposed in [14]: here an inclusion relation is replaced by a cardinality restriction then, when a conflict is detected, the corresponding cardinality restriction is weakened by adapting the number of elements involved. However, when ABox assertions are responsible for the inconsistency, the solution consists in deleting such assertions.

Several works have been proposed in the literature in order to handle inconsistencies in DLs. Most of them try to tackle the more general problem of revising terminologies when two consistent sources K and K' are put together, resulting inconsistent. Due to space limitations, we just mention the most closely related. In [15] the basic idea is to adopt a selection function for choosing among different, consistent subsets of an inconsistent DL knowledge base. Other works [2–7] have addressed the problem of diagnosing incoherent ontologies by identifying a minimal subset of axioms responsible for the inconsistency.

As a main difference with all the above mentioned approaches, in our work we focus on the specific problem of revising a TBox after discovering an exception, described by means of ABox information. We try to tackle the problem of handling exceptions in $\mathcal{ALC}$ knowledge bases by exploiting nonmonotonic Description Logics allowing to represent and reason about *prototypical* properties of concepts, not holding for all the instances of that concept but only for the *typical* (or *most normal*) ones. As far as we know, this is the first attempt to handle exceptions in DLs by exploiting the capabilities of a nonmonotonic DL.

In the above example of (non) flying birds, our approach replaces the inclusion $Bird \sqsubseteq Fly$ with the typicality-based weakened inclusion $\mathbf{T}(Bird) \sqsubseteq Fly$, allowing to express properties holding only for typical birds. In this case, the knowledge base obtained by adding the fact $\neg Fly(tweety)$ (Tweety is a non flying bird) is consistent, and Tweety will be an exceptional bird, preserving all other properties eventually ascribed to birds. In the example of the Nixon diamond, our approach provides the following revised TBox:

$\mathbf{T}(Quacker) \sqsubseteq Pacifist$
$\mathbf{T}(Republican) \sqsubseteq \neg Pacifist$

so that, given the information about Nixon being both a Quaker and a Republican, no conclusion about his position about peace is drawn, since Nixon is both an exceptional Quaker and an exceptional Republican.

3 Description Logics of Typicality

In the recent years, a large amount of work has been done in order to extend the basic formalism of DLs with nonmonotonic reasoning features. The traditional approach is to handle defeasible inheritance by integrating some kind of nonmonotonic reasoning mechanisms [16–20]. A simple but powerful nonmonotonic extension of DLs is proposed in [21]. In this approach, "typical" or "normal" properties can be directly specified by means of a "typicality" operator $\mathbf{T}$ enriching the underlying DL; the typicality operator $\mathbf{T}$ is essentially characterized by the core properties of nonmonotonic reasoning, axiomatized by *preferential logic* **P** in [22].

In this work we refer to the most recent approach proposed in [10], where the authors extend $\mathcal{ALC}$ with $\mathbf{T}$ by considering *rational closure* as defined by Lehman and Magidor [11] for propositional logic. Here the $\mathbf{T}$ operator is intended to enjoy the well-established properties of rational logic **R** . Even if $\mathbf{T}$ is a nonmonotonic operator (so that for instance $\mathbf{T}(Bird) \sqsubseteq Fly$ does not entail that $\mathbf{T}(Bird \sqcap Penguin) \sqsubseteq Fly$), the logic itself is monotonic. Indeed, in this logic it is not possible to monotonically infer from $\mathbf{T}(Bird) \sqsubseteq Fly$, in the absence of information to the contrary, that also $\mathbf{T}(Bird \sqcap Black) \sqsubseteq Fly$. Nor it can be nonmonotonically inferred from $Bird(tweety)$, in the absence of information to the contrary, that $\mathbf{T}(Bird)(tweety)$. Nonmonotonicity is achieved by adapting to $\mathcal{ALC}$ with $\mathbf{T}$ the propositional construction of rational closure. This nonmonotonic extension allows to infer typical subsumptions from the TBox. Intuitively, and similarly to the propositional case, the rational closure construction amounts to assigning a *rank* (a level of exceptionality) to every concept; this rank is used to evaluate typical inclusions of the form $\mathbf{T}(C) \sqsubseteq D$: the inclusion is supported by the rational closure whenever the rank of C is strictly smaller than the rank of $C \sqcap \neg D$. From a semantic point of view, nonmonotonicity is achieved by defining, on the top of $\mathcal{ALC}$ with typicality, a minimal model semantics where the notion of minimality is based on the minimization of the ranks of the domain elements. The problem of extending rational closure to ABox reasoning is also taken into account: in order to ascribe typical properties to individuals, the typicality of an individual is maximized. This is done by minimizing its rank (that is, its level of exceptionality). Let us recall the resulting extension $\mathcal{ALC}^{\mathbf{R}}_{min}\mathbf{T}$ in detail.

Definition 1. *We consider an alphabet of concept names $\mathcal{C}$, of role names $\mathcal{R}$, and of individual constants $\mathcal{O}$. Given $A \in \mathcal{C}$ and $R \in \mathcal{R}$, we define:*

$$C_R := A \mid \top \mid \bot \mid \neg C_R \mid C_R \sqcap C_R \mid C_R \sqcup C_R \mid \forall R.C_R \mid \exists R.C_R$$
$$C_L := C_R \mid \mathbf{T}(C_R)$$

A knowledge base is a pair $(\mathcal{T}, \mathcal{A})$. $\mathcal{T}$ contains a finite set of concept inclusions $C_L \sqsubseteq C_R$. $\mathcal{A}$ contains assertions of the form $C_L(a)$ and $R(a, b)$, where $a, b \in \mathcal{O}$.

We define the semantics of the *monotonic* $\mathcal{ALC} + \mathbf{T_R}$, formulated in terms of rational models: ordinary models of $\mathcal{ALC}$ are equipped with a $preference\ relation$ $<$ on the domain, whose intuitive meaning is to compare the "typicality" of domain elements, that is to say $x < y$ means that x is more typical than y. Typical members of a concept C, that is members of $\mathbf{T}(C)$, are the members x of C that are minimal with respect to this preference relation (s.t. there is no other member of C more typical than x).

Definition 2 (Definition 3 [10]). *A model $\mathcal{M}$ of $\mathcal{ALC} + \mathbf{T_R}$ is any structure $\langle \Delta, <, .^I \rangle$ where: Δ is the domain; $<$ is an irreflexive, transitive and modular (if $x < y$ then either $x < z$ or $z < y$) relation over Δ; $.^I$ is the extension function that maps each concept C to $C^I \subseteq \Delta$, and each role R to $R^I \subseteq \Delta \times \Delta$ in the standard way for $\mathcal{ALC}$ concepts:*

$$\top^I = \Delta$$
$$\bot^I = \emptyset$$
$$(\neg C)^I = \Delta \backslash C^I$$
$$(C \sqcap D)^I = C^I \cap D^I$$
$$(C \sqcup D)^I = C^I \cup D^I$$
$$(\forall R.C)^I = \{x \in \Delta \mid \forall y.(x, y) \in R^I \rightarrow y \in C^I\}$$
$$(\exists R.C)^I = \{x \in \Delta \mid \exists y.(x, y) \in R^I \text{ and } y \in C^I\}$$

whereas for concepts built by means of the typicality operator

$$(\mathbf{T}(C))^I = Min_<(C^I)$$

where $Min_<(S) = \{u : u \in S$ and $\nexists z \in S$ such that $z < u\}$. Furthermore, $<$ satisfies the $Well-Foundedness\ Condition$, i.e., for all $S \subseteq \Delta$, for all $x \in S$, either $x \in Min_<(S)$ or $\exists y \in Min_<(S)$ such that $y < x$.

Definition 3 (Definition 4 [10]). *Given an $\mathcal{ALC} + \mathbf{T_R}$ model $\mathcal{M}$= $\langle \Delta, <, .^I \rangle$, we say that: (i) a model $\mathcal{M}$ satisfies an inclusion $C \sqsubseteq D$ if it holds $C^I \subseteq D^I$; (ii) $\mathcal{M}$ satisfies* *an assertion $C(a)$ if $a^I \in C^I$ and $\mathcal{M}$ satisfies an assertion $R(a,b)$ if $(a^I, b^I) \in R^I$. Given a knowledge base K=(TBox,ABox), we say that: (i) $\mathcal{M}$ satisfies TBox if $\mathcal{M}$ satisfies all inclusions in TBox; (ii) $\mathcal{M}$ satisfies ABox if $\mathcal{M}$ satisfies all assertions in ABox; (iii) $\mathcal{M}$ satisfies K if it satisfies both its TBox and its ABox.*

Given a knowledge base K, an inclusion $C_L \sqsubseteq C_R$, and an assertion $C_L(a)$, with $a \in \mathcal{O}$, we say that the inclusion $C_L \sqsubseteq C_R$ is derivable from K, written $K \models_{\mathcal{ALC}^{\mathbf{R}}\mathbf{T}} C_L \sqsubseteq C_R$, if ${C_L}^I \subseteq {C_R}^I$ holds in all models $\mathcal{M} = \langle \Delta, <, .^I \rangle$ satisfying K. Moreover, we say the assertion $C_L(a)$ is derivable from K, written $K \models_{\mathcal{ALC}^{\mathbf{R}}\mathbf{T}} C_L(a)$, if $a^I \in {C_L}^I$ holds in all models $\mathcal{M} = \langle \Delta, <, .^I \rangle$ satisfying K.

Definition 4 (Rank of a domain element $k_{\mathcal{M}}(x)$, Definition 5 [10]). *Given a model $\mathcal{M} = \langle \Delta, <, .^I \rangle$, the rank $k_{\mathcal{M}}$ of a domain element $x \in \Delta$, is the length of the longest chain $x_0 < \ldots < x$ from x to a minimal x_0 (i.e. such that there is no x' such that $x' < x_0$).*

As already mentioned, although the typicality operator $\mathbf{T}$ itself is nonmonotonic (i.e. $\mathbf{T}(C) \sqsubseteq D$ does not imply $\mathbf{T}(C \sqcap E) \sqsubseteq D$), the logic $\mathcal{ALC} + \mathbf{T_R}$ is monotonic: what is inferred from K can still be inferred from any K' with $K \subseteq K'$. This is a clear limitation in DLs. As a consequence of the monotonicity of $\mathcal{ALC} + \mathbf{T_R}$, one cannot deal with irrelevance. For instance, from the knowledge base of birds and penguins, one cannot derive that $K \models_{\mathcal{ALC}^{\mathbf{R}}\mathbf{T}} \mathbf{T}(Penguin \sqcap Black) \sqsubseteq \neg Fly$, even if the property of being black is irrelevant with respect to flying. In the same way, if we added to K the information that Tweety is a bird ($Bird(tweety)$), in $\mathcal{ALC} + \mathbf{T_R}$ one cannot tentatively derive, in the absence of information to the contrary, that $\mathbf{T}(Bird)(tweety)$ and $Fly(tweety)$.

In order to tackle this problem, in [10] the definition of rational closure introduced by Lehmann and Magidor [11] for the propositional case has been extended to the DL $\mathcal{ALC} + \mathbf{T_R}$. The resulting nonmonotonic logic is called $\mathcal{ALC}^{\mathbf{R}}_{min}\mathbf{T}$. From a semantic point of view, in [10] it is shown that minimal rational models that minimize the rank of domain elements can be used to give a semantical reconstruction of this extension of rational closure. The idea is as follows: given two models of K, one in which a given domain element x has rank x_1 and another in which it has rank x_2, with $x_1 > x_2$, then the latter is preferred, as in this model the element x is "more normal" than in the former.

Definition 5 (Minimal models, Definition 8 [10]). *Given $\mathcal{M} = \langle \Delta, <, .^I \rangle$ and $\mathcal{M}' = \langle \Delta', <', .^{I'} \rangle$ we say that $\mathcal{M}$ is preferred to $\mathcal{M}'$ ($\mathcal{M} <_{FIMS} \mathcal{M}'$) if: (i) $\Delta = \Delta'$; (ii) $C^I = C^{I'}$ for all concepts C; (iii) for all $x \in \Delta$, it holds that $k_{\mathcal{M}}(x) \leq k_{\mathcal{M}'}(x)$ whereas there exists $y \in \Delta$ such that $k_{\mathcal{M}}(y) < k_{\mathcal{M}'}(y)$. Given a knowledge base K=(TBox,ABox), we say that $\mathcal{M}$ is a minimal model of K with respect to $<_{FIMS}$ if it is a model satisfying K and there is no $\mathcal{M}'$ model satisfying K such that $\mathcal{M}' <_{FIMS} \mathcal{M}$. Furthermore, we say that $\mathcal{M}$ is preferred to $\mathcal{M}'$ with respect to ABox, and we write $\mathcal{M} <_{ABox} \mathcal{M}'$, if, for all individual constants a occurring in ABox, it holds that $k_{\mathcal{M}}(a^I) \leq k_{\mathcal{M}'}(a^{I'})$ and there is at least one individual constant b occurring in ABox such that $k_{\mathcal{M}}(b^I) < k_{\mathcal{M}'}(b^{I'})$.*

Given a knowledge base $K = (\mathcal{T}, \mathcal{A})$, in [10] it is shown that an inclusion $C \sqsubseteq D$ (respectively, an assertion $C(a)$) belongs to the rational closure of K if and only if $C \sqsubseteq D$ (resp., $C(a)$) holds in all minimal models of K of a "special" kind, named *canonical models*. The rational closure construction for $\mathcal{ALC}$ is inexpensive, since it retains the same complexity of the underlying logic, and thus a good candidate to define effective nonmonotonic extensions of DLs. More precisely, the problem of deciding whether a typical inclusion belongs to the rational closure of $\mathcal{T}$ is in EXPTIME as well as the problem of deciding whether an assertion $C(a)$ belongs to the rational closure over $\mathcal{A}$.

4 Typicality-based revision

Similarly to what done in [9], we define a notion of revised knowledge base, precisely a *typicality-based* revised knowledge base. Given a consistent knowledge base

$$K = (\mathcal{T}, \mathcal{A})$$

we have to tackle the problem of accommodating a further ABox information $\mathcal{A}'$, describing an individual x that belongs to the extensions of the concepts $D_1, D_2, \ldots, D_n$, and that is an *exception* to the knowledge described by K. Given a consistent ABox

$$\mathcal{A}' = \{D_1(x), D_2(x), \ldots, D_n(x)\}$$

we have that the knowledge base

$$(\mathcal{T}, \mathcal{A} \cup \mathcal{A}')$$

is inconsistent. We revise the TBox $\mathcal{T}$ of K by replacing some standard inclusions $C \sqsubseteq D$ with typicality inclusions $\mathbf{T}(C) \sqsubseteq D$, in a way such that the resulting revised knowledge base $(\mathcal{T}^{new}, \mathcal{A} \cup \mathcal{A}')$ is consistent in $\mathcal{ALC}^{\mathbf{R}}_{min}\mathbf{T}$. We first define a typicality-based weakening of a TBox inclusion as follows:

Definition 6 (Typicality-based weakening of $C \sqsubseteq D$). *Given an $\mathcal{ALC}$ inclusion $C \sqsubseteq D$, we say that either $C \sqsubseteq D$ or $\mathbf{T}(C) \sqsubseteq D$ is a typicality-based weakening of $C \sqsubseteq D$.*

Definition 7 (Typicality-based weakened TBox). *Let $K = (\mathcal{T}, \mathcal{A})$ be a consistent knowledge base and $\mathcal{A}' = \{D_1(x), D_2(x), \ldots, D_n(x)\}$ be a consistent ABox, and suppose that $(\mathcal{T}, \mathcal{A} \cup \mathcal{A}')$ is inconsistent. An $\mathcal{ALC}^{\mathbf{R}}_{min}\mathbf{T}$ TBox $\mathcal{T}_i$ is a typicality-based weakening of $\mathcal{T}$ if the following conditions hold:*

- *$(\mathcal{T}_i, \mathcal{A} \cup \mathcal{A}')$ is consistent*
- *there exists a bijection f from $\mathcal{T}$ to $\mathcal{T}_i$ such that for each $C \sqsubseteq D \in \mathcal{T}$, $f(C \sqsubseteq D)$ is a typicality-based weakening of $C \sqsubseteq D$.*

Given $K = (\mathcal{T}, \mathcal{A})$, we denote with $Rev_{\mathbf{T},\mathcal{A}'}(K)$ the set of all typicality-based weakenings of $\mathcal{T}$ given a newly received ABox $\mathcal{A}'$.

Among all typicality-based weakenings in $Rev_{\mathbf{T},\mathcal{A}'}(K)$ we select the one, called $\mathcal{T}^{new}$, capturing a notion of *minimal changes* needed in order to accommodate the discovered exception. The resulting $\mathcal{ALC}^{\mathbf{R}}_{min}\mathbf{T}$ knowledge base

$$K^{new} = (\mathcal{T}^{new}, \mathcal{A} \cup \mathcal{A}')$$

is consistent.

Let us introduce the typicality-based revision approach by means of some examples.

Example 1. Let $K = (\mathcal{T}, \emptyset)$ where $\mathcal{T}$ is:

$$\begin{array}{ll} \mathit{Quacker} \sqsubseteq \mathit{Christian} & (1)\\ \mathit{Christian} \sqsubseteq \mathit{Pacifist} & (2)\\ \mathit{RepublicanPresident} \sqsubseteq \mathit{Republican} & (3)\\ \mathit{Republican} \sqsubseteq \neg \mathit{Pacifist} & (4) \end{array}$$

and let $\mathcal{A}' = \{\mathit{Quacker}(\mathit{nixon}), \mathit{RepublicanPresident}(\mathit{nixon})\}$.

It can be shown that there are eleven different typicality-based weakenings in $Rev_{\mathbf{T},\mathcal{A}'}(K)$, but the one chosen as the typicality-based revision $\mathcal{T}^{new}$ considers Nixon as an exceptional Christian and Republican, as follows:

$$\begin{array}{l} \mathit{Quacker} \sqsubseteq \mathit{Christian}\\ \mathbf{T}(\mathit{Christian}) \sqsubseteq \mathit{Pacifist}\\ \mathit{RepublicanPresident} \sqsubseteq \mathit{Republican}\\ \mathbf{T}(\mathit{Republican}) \sqsubseteq \neg \mathit{Pacifist} \end{array}$$

$\mathcal{T}^{new}$ is a weakening of K given $\mathcal{A}'$, therefore it accommodates the exception represented by the individual x. The chosen solution does not lead to trivially inconsistent concepts: in Example 1 above, let us consider a revised TBox in which only the inclusion (4) is weakened; that is, $\mathbf{T}(\mathit{Republican}) \sqsubseteq \neg \mathit{Pacifist}$ is the only typicality inclusion introduced. Of course, such a TBox is a weakening of K, and the resulting knowledge base is consistent; however, all its models are such that there are not typical Republicans being also Quakers, in other words the concept $\mathbf{T}(\mathit{Republican}) \sqcap \mathit{Quacker}$ is inconsistent. As a consequence, Nixon is an exceptional Republican (i.e., a non-typical Republican); however Nixon is also a Quaker, and then a Christian and, therefore, it can be inferred that he is pacifist. This is, in general, an unwanted conclusion, at least when reasoning in a skeptical way.

Moreover, the discovered exception is not handled in a "lazy way", by weakening all the inclusion relations in the TBox. As a matter of fact, when an inclusion $C \sqsubseteq D$ is weakened by $\mathbf{T}(C) \sqsubseteq D$, we are moving from models in $\mathcal{ALC}$ where all Cs are Ds ($C^I \subseteq D^I$) to models in $\mathcal{ALC}^{\mathbf{R}}_{min}\mathbf{T}$ where we are able to distinguish between typical Cs and exceptional ones (C^I and $\mathbf{T}(C)^I$), and x is an exceptional C (i.e., $x \not\in \mathbf{T}(C)^I$). However, we prefer a knowledge base whose models *minimize* the number of concepts for which x is exceptional. For instance, the following TBox is a weakening of K in Example 1:

$$\begin{array}{l} \mathbf{T}(\mathit{Quacker}) \sqsubseteq \mathit{Christian}\\ \mathbf{T}(\mathit{Christian}) \sqsubseteq \mathit{Pacifist}\\ \mathbf{T}(\mathit{RepublicanPresident}) \sqsubseteq \mathit{Republican}\\ \mathbf{T}(\mathit{Republican}) \sqsubseteq \neg \mathit{Pacifist} \end{array}$$

However, in this case we are introducing the–not needed–opportunity of having exceptional Quakers not being Christians, as well as the counterintuitive fact of having exceptional Republican presidents not being Republicans.

Finally, our approach prefers a revision in which the typicality-based weakenings are performed on the most general concepts. Consider again Example 1: obviously, also the following

$\mathbf{T}(Quacker) \sqsubseteq Christian$
$Christian \sqsubseteq Pacifist$
$\mathbf{T}(RepublicanPresident) \sqsubseteq Republican$
$Republican \sqsubseteq \neg Pacifist$

is a weakening of K, however this corresponds to consider Nixon as an exceptional Republican president, not allowing to conclude that he is Republican ($Republican(nixon)$ is not entailed by the revised knowledge base). Symmetrically, Nixon is considered as an exceptional Quaker, not allowing to infer that he is a Christian (again $Christian(nixon)$ is not entailed by the revised knowledge base). We prefer $\mathcal{T}^{new}$ since the typicality operator is introduced over the most general concepts $Christian$ and $Republican$, and as a consequence, from the revised knowledge base, one can infer that Nixon is a Christian and that he is Republican, but (correctly) no conclusions about being pacifist or not are drawn (Nixon is both an exceptional Christian and an exceptional Republican).

5 Computing a Revised TBox

In this section we propose a possible algorithm that revises a given knowledge base $K = (\mathcal{T}, \mathcal{A})$ according to the typicality-based weakening described in the previous section. As describe above, we are interested in revising K in order to accommodate an additional observation of the form $\mathcal{A}' = \{D_1(x), \dots, D_n(x)\}$ such that $K = (\mathcal{T}, \mathcal{A} \cup \mathcal{A}')$ is inconsistent. Intuitively, the individual x is exceptional as it possesses properties that, as far as we know in the current K, do not occur together, but rather are in contrast.

The algorithm we propose relies on two main concepts: the computation of a Minimal Unsatisfiability-Preserving Sub-TBoxes (*mups*), that singles out the subset of inclusions strictly involved in the inconsistency; and the notion of Generalized Qualified Subconcepts (*gqs*). Both concepts have been introduced by Schlobach et al. in their seminal works [2, 3] about the problem of debugging a DL terminology. We shortly recall these basic notions before presenting our algorithm. Note that the techniques proposed in [2, 3] are restricted to *unfoldable* TBoxes, only containing unique, acyclic definitions. An axiom is called a definition of A if it is of the form $A \sqsubseteq C$, where $A \in \mathcal{C}$ is an atomic concept. An axiom $A \sqsubseteq C$ is unique if the KB contains no other definition of A. An axiom is acyclic if C does not refer either directly or indirectly (via other axioms) to A [1]. Since we rely on the above mentioned works, from now on we restrict our concern to unfoldable TBoxes.

5.1 Mups, gqs, and specificity ordering

To explain incoherences in terminologies, Schlobach et al. propose a methodology based on two steps: first, *axiom pinpointing* excludes axioms which are irrelevant to the incoherence; second, *concept pinpointing* provides a simplified definition highlighting the exact position of a contradiction within the axioms previously selected. In this paper we are interested in the axiom pinpointing step, which identifies debugging-relevant axioms. Intuitively, an axiom is relevant for debugging if, when removed, a TBox becomes consistent, or at least one previously unsatisfiable concept turns satisfiable. The notion of subset of relevant axioms is captured by the following definition.

Definition 8 (MUPS, Definition 3.1 [3]). *Let C be a concept which is unsatisfiable in a TBox $\mathcal{T}$. A set $\mathcal{T}' \subseteq \mathcal{T}$ is a minimal unsatisfiability-preserving sub-TBox (MUPS) of $\mathcal{T}$ if C is unsatisfiable in $\mathcal{T}'$, and C is satisfiable in every sub-TBox $\mathcal{T}'' \subset \mathcal{T}'$.*

In the following, $mups(\mathcal{T}, C)$ is used to denote the set of MUPS for a given terminology $\mathcal{T}$ and a concept C. Intuitively, each set of axioms in $mups(\mathcal{T}, C)$ represents a conflict set; i.e., a set of axioms that cannot all be satisfied. From this point of view, it is therefore possible to infer a diagnosis for the concept C by applying the Hitting-Set tree algorithm proposed by Reiter [23]. However, the set $mups(\mathcal{T}, C)$ is sufficient for our purpose of dealing with exceptions.

In [2], Schlobach also proposes a methodology for explaining concept subsumptions. The idea is to reduce the structural complexity of the original concepts in order to highlight the logical interplay between them. To this aim, Schlobach proposes to exploit the structural similarity of concepts, that can be used to simplify terminological concepts, and hence the subsumption relations. The structural similarity is based on the notion of *qualified subconcepts*; namely, variants of those concepts a knowledge engineer explicitly uses in the modeling process, and where the context (i.e., sequence of quantifiers and number of negations) of this use is kept intact. Schlobach specifies the notion of qualified subconcepts in two ways: Generalized Qualified Subconcepts (gqs), and Specialized Qualified Subconcepts (sqs) which are defined by induction as follows:

Definition 9 (Generalized/Specialized Qualified Subconcepts, [2]). *Given concepts A, C and D, we define:*

$gqs(A) = sqs(A) = \{A\}$ *if A is atomic*
$gqs(C \sqcap D) = \{C', D', C' \sqcap D' | C' \in gqs(C), D' \in gqs(D)\}$
$gqs(C \sqcup D) = \{C' \sqcup D' | C' \in gqs(C), D' \in gqs(D)\}$
$gqs(\exists r.C) = \{\exists r.C' | C' \in gqs(C)\}$
$gqs(\forall r.C) = \{\forall r.C' | C' \in gqs(C)\}$
$gqs(\neg C) = \{\neg C' | C' \in sqs(C)\}$
$sqs(C \sqcap D) = \{C' \sqcap D' | C' \in sqs(C), D' \in sqs(D)\}$
$sqs(C \sqcup D) = \{C', D', C' \sqcup D' | C' \in sqs(C), D' \in sqs(D)\}$
$sqs(\exists r.C) = \{\exists r.C' | C' \in sqs(C)\}$
$sqs(\forall r.C) = \{\forall r.C' | C' \in sqs(C)\}$
$sqs(\neg C) = \{\neg C' | C' \in gqs(C)\}$

As Schlobach himself notes, a simple consequence of this definition is that $\models C \sqsubseteq C'$ for every $C' \in gqs(C)$, and $\models D' \sqsubseteq D$ for each $D' \sqsubseteq sqs(D)$.

We slightly extend the base case of Definition 9 as follows:

Definition 10 (Extended GQS and SQS). *Given a TBox $\mathcal{T}$, we define $sqs(C)$ and $gqs(C)$ by adding the following clauses to those in Definition 9:*

$gqs(A) = \{A\} \cup \{gqs(D) \mid A \sqsubseteq D \in \mathcal{T}\}$ *if A is atomic*
$sqs(A) = \{A\} \cup \{sqs(C) \mid C \sqsubseteq A \in \mathcal{T}\}$ *if A is atomic*
$gqs(\neg A) = \{\neg A\} \cup \{gqs(D) \mid \neg A \sqsubseteq D \in \mathcal{T}\}$ *if A is atomic*
$sqs(\neg A) = \{\neg A\} \cup \{sqs(C) \mid C \sqsubseteq \neg A \in \mathcal{T}\}$ *if A is atomic*

Thus, we also take into account the axioms (i.e., concept inclusions) defined in a given TBox. This generalization allows us to move upward (by means of gqs), and downward (by means of sqs) in the hierarchy of concepts defined by a given TBox $\mathcal{T}$. Relying on the notions of (extended) sqs and gqs, we can define a partial ordering relation between concepts as follows:

Definition 11. *Let C and D be two concepts ($C \neq D$) in a given* **TBox** *$\mathcal{T}$, we say that C is more specific than D, denoted as $C \prec D$, iff at least one of the following relations holds: (i) $C \in sqs(D)$, or (ii) $D \in gqs(C)$.*

It is easy to see that $\prec$ is irreflexive, antisymmetric, and transitive; however, it is just partial because $\prec$ is not defined for any pair of concepts; i.e., there may exist two concepts C and D such that neither $C \prec D$ nor $D \prec C$ holds. As we will discuss later, the methodology we propose for determining which concepts represent properties to be made "typical" relies on the fact that concepts are considered in order from the most specific to the most general. In those situations where two concepts are not directly comparable with one another by means of $\prec$, either ordering is possible.

By extension, we can order two concept inclusions $C \sqsubseteq D' \prec C \sqsubseteq D''$, meaning that $C \sqsubseteq D'$ is more specific than $C \sqsubseteq D''$, as far as $D' \prec D''$.

5.2 Algorithm

We are now in the position for presenting an algorithmic solution to the problem of accommodating an exceptional individual x within an existing knowledge base $K = (\mathcal{T}, \mathcal{A})$. A central role is played by the notion of *mups* since it allows us to isolate a portion of $\mathcal{T}$ which is minimal and relevant. In the original formulation by Schlobach et al., however, a *mups* is computed given a concept which is unsatisfiable in a given terminology. Since in our problem $\mathcal{T}$ is consistent, we need to add in $\mathcal{T}$ a novel concept inclusion which, on the one side represents the novel observation about x, and on the other side represents an unsatisfiable concept from which the computation of a *mups* can start. We therefore create a temporary terminology $\mathcal{T}' = \mathcal{T} \cup \{X \sqsubseteq D_1 \sqcap \ldots \sqcap D_n\}$, where the fake concept X simply reflects the properties of the exceptional individual x. The $mups(\mathcal{T}', X)$ can now be computed following the methodology suggested in [3]. Note that in principle $mups(\mathcal{T}', X)$ is a set of alternative sub-TBoxes, in this paper, however, we assume that $mups(\mathcal{T}', X)$ only contains a single sub-TBox, and leave as future work the study of how handling *mups* containing alternative sub-TBoxes that possibly lead to different revised knowledge bases.

As noted above, we are interested in revising K so that it satisfies a minimality criterion about the exceptionality of x. Algorithmically, this requires to consider all the possible generalizations of X, computed according to the extended notions of *gqs* and *sqs* given in Definition 10. This structure is called $\mathcal{GQS}(X)$ and is defined as follows:

Definition 12 **($\mathcal{GQS}(X)$).** *Let $\mathcal{T}'$ be the terminology $\mathcal{T}$ modified with the addition of the concept inclusion describing the exceptional, observed individual x, i.e. $\mathcal{T}' = \mathcal{T} \cup \{X \sqsubseteq D_1 \sqcap \ldots \sqcap D_n\}$, and let $mups(\mathcal{T}', X)$ be the minimal subset of $\mathcal{T}$ preserving the inconsistency with X; the $\mathcal{GQS}(X)$ is a list of concept inclusions, ordered according to the specificity ordering $\prec$ in Definition 11, such that each element $X \sqsubseteq E_1 \sqcap \ldots \sqcap E_m$ in this list satisfies the following conditions:*

1. *each $E_1, \ldots, E_m$ is a concept mentioned in $mups(\mathcal{T}', X)$;*
2. *$E_1 \sqcap \ldots \sqcap E_m$ is not trivially inconsistent (the same concept does not appear both with and without a negation), i.e. there are not i, j such that $E_i = C$ and $E_j = \neg C$;*
3. *$E_1 \sqcap \ldots \sqcap E_m$ does not contradict what we know about X from the direct observation of x (e.g., if $X \sqsubseteq D$, then $\neg D$ cannot occur in place of any $E_1, \ldots, E_m$).*

Condition 1 puts a limit within which the generalizations of X have to be looked for. The *mups*, in fact, restricts the concepts of interest to only those concepts that are strictly involved in the inconsistency of X. Condition 2 discards any inclusion that is not informative at all because it is inconsistent; observe that inconsistent inclusions can be generated because the starting terminology $\mathcal{T}'$ is inconsistent itself. Condition 3 discards any inclusion that contradicts the observations about x, which must always be considered as correct. Note that the computation of $\mathcal{GQS}(X)$ obtained as a simple rewriting of the concepts $D_1, \ldots, D_n$ in terms of their respective *gqs* always terminates since, on the one hand, we assume that $\mathcal{T}$ is acyclic and unfoldable and, on the other hand, the set of inclusions and concepts to be considered is limited to the ones in $mups(\mathcal{T}', X)$.

Since we adopt a skeptical approach, the intuition is that the revised knowledge base $K^{new} = (\mathcal{T}^{new}, \mathcal{A} \cup \mathcal{A}')$ must be pairwise consistent with each of the inclusions in $\mathcal{GQS}(X)$; that is, K^{new} must be ready to accommodate any further observations about x, and hence it must not make assertions about x which have not been directly observed. For this reason, the revised $\mathcal{T}^{new}$ must differ from the original $\mathcal{T}$ only for a finite number of inclusions weakened by the operator $\mathbf{T}$.

Algorithm 1 Revise ($\mathcal{T}$, $X \sqsubseteq D_1 \sqcap \ldots \sqcap D_n$)

$\mathcal{T}' \leftarrow \mathcal{T} \cup \{X \sqsubseteq D_1 \sqcap \ldots \sqcap D_n\}$
$\mathcal{T}^{mups} \leftarrow mups(\mathcal{T}', X)$
$\mathcal{T}^{rev} \leftarrow \mathcal{T}^{mups} \setminus \{X \sqsubseteq D_1 \sqcap \ldots \sqcap D_n\}$
Compute $\mathcal{GQS}(X)$
for all inclusions $X \sqsubseteq E_1 \sqcap \ldots \sqcap E_m$ in $\mathcal{GQS}(X)$ from the most general inclusion to the most specific **do**
  for all inclusions $C \sqsubseteq C_1 \sqcap \ldots \sqcap C_k$ in $\mathcal{T}^{rev}$ **do**
    if there exist E_i and C_j such that $E_i = F$ and $C_j = \neg F$ **then**
      $\mathcal{T}^{rev} \leftarrow \mathcal{T}^{rev} \setminus \{C \sqsubseteq C_1 \sqcap \ldots \sqcap C_k\} \cup \{\mathbf{T}(C) \sqsubseteq C_1 \sqcap \ldots \sqcap C_k\}$
$\mathcal{T}^{new} \leftarrow$ update $\mathcal{T}$ with weakened inclusions according to $\mathcal{T}^{rev}$
return $\mathcal{T}^{new}$

The algorithm for revising $K = (\mathcal{T}, \mathcal{A})$ into K^{new} is given in Algorithm 1. The algorithm takes as inputs the original terminology $\mathcal{T}$ and the fake inclusion $X \sqsubseteq D_1 \sqcap \ldots \sqcap D_n$. The first step consists in computing $mups(\mathcal{T}', X)$; the resulting sub-terminology, denoted as $\mathcal{T}^{mups}$ contains, besides $X \sqsubseteq D_1 \sqcap \ldots \sqcap D_n$, a minimal set of inclusions in $\mathcal{T}$ which are involved in the inconsistency caused by the exceptional individual x. Since we are interested in revising the inclusions in the original terminology, we isolate these concepts in the sub-terminology $\mathcal{T}^{rev}$. As noted above, we assume that $mups(\mathcal{T}', X)$ results in a single sub-TBox; in case alternative sub-TBoxes there existed, we assume to select one of them; e.g., the one with the least number of inclusions could be heuristically selected.

The next step is the computation of $\mathcal{GQS}(X)$. After these two preliminary steps, the algorithm looks for the concepts in $\mathcal{T}^{rev}$ to be weakened. The pinpointing of these concepts proceeds as a pairwise consistency check between $\mathcal{T}^{rev}$ and the concept inclusions in $\mathcal{GQS}(X)$, considered from the most general to the most specific. In this way the algorithm answers to two preference requirements over the revised K^{new}: (1) the inclusions to be weakened must be the most general possible with respect to X, and (2) the number of weakened inclusions must be minimal, matching the desiderata for the typicality-based revised TBox described in the previous section. In fact, by considering the inclusions in $\mathcal{GQS}(X)$ from the most general, the algorithm finds first those inclusions in $\mathcal{T}^{rev}$ mentioning concepts that are the greatest generalizations of X, and hence represent the most general properties of the individual x that make x exceptional. On the other hand, the weakening of an inclusion $C \sqsubseteq C_1 \sqcap \ldots \sqcap C_k \in \mathcal{T}^{rev}$ makes consistent any other concept that belongs to $sqs(C)$, and this guarantees a minimal number of weakenings. The algorithm terminates with the generation of $\mathcal{T}^{new}$, which is simply the original $\mathcal{T}$ updated according to the weakened concepts found in the working terminology $\mathcal{T}^{rev}$ in line 9.

It can be shown that the knowledge base $K^{new} = (\mathcal{T}^{new}, \mathcal{A}^{new})$ where $\mathcal{A}^{new} = \mathcal{A} \cup \mathcal{A}'$ and $\mathcal{T}^{new}$ results from algorithm *Revise* in Algorithm 1, is consistent.

Let us conclude this section with an example: a simplified variant of the Pizza ontology distributed together with Protégé.

Example 2. Let us consider $K = (\mathcal{T}_{tops}, \emptyset)$, where $\mathcal{T}_{tops}$ contains the following subsumption relations:

$ax_1 : \mathit{FatToppings} \sqsubseteq \mathit{PizzaToppings}$,
$ax_2 : \mathit{CheesyToppings} \sqsubseteq \mathit{FatToppings}$,
$ax_3 : \mathit{VegetarianToppings} \sqsubseteq \neg \mathit{FatToppings}$.

Now, let us assume that we discover that there also exists the *tofu* cheese, which is made of curdled soybeanmilk. Formally, we have a newly received ABox

$$\mathcal{A}' = \{\mathit{VegetarianToppings}(\mathit{tofu}), \mathit{CheesyToppings}(\mathit{tofu})\}.$$

Obviously, the knowledge base $K = (\mathcal{T}_{tops}, \mathcal{A}')$ is inconsistent. We apply the methodology introduced in this work in order to find the typicality-based revision of the TBox $\mathcal{T}_{tops}$: in this case, among the typicality-based weakenings of $\mathcal{T}_{tops}$, we aim at obtaining the following $\mathcal{T}_{tops}^{new}$:

$\mathit{FatToppings} \sqsubseteq \mathit{PizzaToppings}$,
$\mathbf{T}(\mathit{CheesyToppings}) \sqsubseteq \mathit{FatToppings}$,
$\mathbf{T}(\mathit{VegetarianToppings}) \sqsubseteq \neg \mathit{FatToppings}$

Let us see how the algorithm finds such a revised terminology. The first step consists in constructing $\mathcal{T}'_{tops}$ by adding the following inclusion to $\mathcal{T}_{tops}$:

$$ax_4 : X \sqsubseteq \mathit{CheesyToppings} \sqcap \mathit{VegetarianToppings}.$$

Then we can compute $\mathcal{T}^{mups}$ =$\{ax_2, ax_3, ax_4\}$; it follows that $\mathcal{T}_{tops}^{rev} = \{ax_2, ax_3\}$. The algorithm proceeds by computing $\mathcal{GQS}(X) = \{\ g_1 : X \sqsubseteq CheesyToppings \sqcap VegetarianToppings, g_2 : X \sqsubseteq CheesyToppings, g_3 : X \sqsubseteq VegetarianToppings, g_4 : X \sqsubseteq FatToppings \sqcap VegetarianToppings, g_5 : X \sqsubseteq CheesyToppings \sqcap$

$\neg FatToppings$, $g_6 : X \sqsubseteq FatToppings$, $g_7 : X \sqsubseteq \neg FatToppings$ }. Note that the inclusion $X \sqsubseteq FatToppings \sqcap \neg FatToppings$ is discarded as it is trivially contradictory. $\mathcal{GQS}(X)$ thus keeps a spectrum of possible generalizations of X that are not trivial and consistent with the observation we have about *tofu* (i.e., being both *CheesyToppings* and *VegetarianToppings*). As such, $\mathcal{T}^{rev}_{tops}$ must be consistent with all of them. The algorithm, thus, considers all the inclusions in $\mathcal{GQS}(X)$ from the tail of the list (i.e., from the most general inclusions), and verifies, in a pairwise way, the consistency of each of them with ax_2 and ax_3. It is easy to see that, while g_7 and g_6 are pairwise consistent with ax_2 and ax_3, g_5 is inconsistent with ax_2, and g_4 is inconsistent with ax_3. The two inclusions ax_2 and ax_3 are therefore weakened obtaining the expected revised terminology $\mathcal{T}^{new}_{tops}$.

6 Conclusions and Future Issues

We have presented a typicality-based revision of a DL TBox in presence of exceptions. We exploit techniques and algorithms proposed in [2, 3], which have been extended to more expressive DLs such as $\mathcal{SHOIN}$ in [5], corresponding to ontology language OWL-DL. We aim at extending our typicality-based revision to such expressive DLs in future research, by exploiting results provided in [24, 25] where the typicality operator and the rational closure construction have been applied to the logics $\mathcal{SHIQ}$ and $s\mathcal{ROIQ}$.

As mentioned in Section 5, a drawback of the debugging approach in [2, 3] is that it is restricted to *unfoldable* TBoxes, only containing unique, acyclic definitions. This restriction could seem too strong for our objective of representing and reasoning about defeasible inheritance in a natural way. As an example, a TBox expressing that students are not tax payers, but working students do pay taxes, could be naturally expressed by the following, not unfoldable, TBox=$\{Student \sqsubseteq \neg TaxPayer, Student \sqcap Worker \sqsubseteq TaxPayer\}$. In order to fill this gap, in [7, 5] axiom pinpointing is extended to general TBoxes. A set of algorithms for computing axiom pinpointing, in particular to compute the set of MUPS for a given terminology $\mathcal{T}$ and a concept A, is also provided. Another aspect that deserves further investigation is the extension of our approach to revise not unfoldable TBoxes. Furthermore, we intend to develop an implementation of the proposed algorithms, by considering the integration with existing tools for manually modifying ontologies when inconsistencies are detected.

References

1. Baader, F., Calvanese, D., McGuinness, D., Nardi, D., Patel-Schneider, P.: The Description Logic Handbook - Theory, Implementation, and Applications. Cambridge (2007)
2. Schlobach, S., Cornet, R.: Explanation of terminological reasoning: A preliminary report. In: Description Logics. (2003)
3. Schlobach, S., Cornet, R.: Non-standard reasoning services for the debugging of description logic terminologies. In: IJCAI. (2003) 355–362
4. Schlobach, S.: Diagnosing terminologies. In: AAAI. (2005) 670–675
5. Kalyanpur, A., Parsia, B., Cuenca-Grau, B., Sirin, E.: Axiom pinpointing: Finding (precise) justifications for arbitrary entailments in $\mathcal{SHOIN}$ (owl-dl). Technical report, UMIACS, 2005-66 (2006)

6. Kalyanpur, A., Parsia, B., Sirin, E., Grau, B.C.: Repairing unsatisfiable concepts in owl ontologies. In: ESWC. (2006) 170–184
7. Parsia, B., Sirin, E., Kalyanpur, A.: Debugging owl ontologies. In: WWW. (2005) 633–640
8. Micalizio, R., Pozzato, G.L.: Revising description logic terminologies to handle exceptions: a first step. In Giordano, L., Gliozzi, V., Pozzato, G.L., eds.: Proceedings of the 29th Italian Conference on Computational Logic, Torino, Italy, June 16-18, 2014. Volume 1195 of CEUR Workshop Proceedings., CEUR-WS.org (2014) 225–240
9. Qi, G., Liu, W., Bell, D.A.: A revision-based approach to handling inconsistency in description logics. Artificial Intelligence Review **26**(1-2) (2006) 115–128
10. Giordano, L., Gliozzi, V., Olivetti, N., Pozzato, G.L.: Semantic characterization of rational closure: From propositional logic to description logics. Artif. Intelligence **226** (2015) 1–33
11. Lehmann, D., Magidor, M.: What does a conditional knowledge base entail? Artificial Intelligence **55**(1) (1992) 1–60
12. Benferhat, S., Kaci, S., Berre, D.L., Williams, M.: Weakening conflicting information for iterated revision and knowledge integration. Artif. Intell. **153**(1-2) (2004) 339–371
13. Benferhat, S., Baida, R.E.: A stratified first order logic approach for access control. International Journal of Intelligent Systems **19**(9) (2004) 817–836
14. Meyer, T.A., Lee, K., Booth, R.: Knowledge integration for description logics. In Veloso, M.M., Kambhampati, S., eds.: Proceedings, The 20th National Conference on Artificial Intelligence and the 17th Innovative Applications of Artificial Intelligence Conference, July 9-13, 2005, Pittsburgh, Pennsylvania, USA, AAAI Press / The MIT Press (2005) 645–650
15. Huang, Z., van Harmelen, F., ten Teije, A.: Reasoning with inconsistent ontologies. In Kaelbling, L.P., Saffiotti, A., eds.: IJCAI-05, Proc. of the 19th International Joint Conference on Artificial Intelligence, Edinburgh, Scotland, UK, July 30-August 5, 2005. (2005) 454–459
16. Giordano, L., Gliozzi, V., Olivetti, N., Pozzato, G.L.: ALC+T: a preferential extension of Description Logics. Fundamenta Informaticae **96** (2009) 1–32
17. Baader, F., Hollunder, B.: Embedding defaults into terminological knowledge representation formalisms. J. Autom. Reasoning **14**(1) (1995) 149–180
18. Ke, P., Sattler, U.: Next Steps for Description Logics of Minimal Knowledge and Negation as Failure. In Baader, F., Lutz, C., Motik, B., eds.: Proceedings of Description Logics. Volume 353 of CEUR Workshop Proceedings., Dresden, Germany, CEUR-WS.org (May 2008)
19. Casini, G., Straccia, U.: Rational Closure for Defeasible Description Logics. In Janhunen, T., Niemelä, I., eds.: Proceedings of the 12th European Conference on Logics in Artificial Intelligence (JELIA 2010). Volume 6341 of Lecture Notes in Artificial Intelligence., Helsinki, Finland, Springer (September 2010) 77–90
20. Casini, G., Straccia, U.: Defeasible Inheritance-Based Description Logics. In Walsh, T., ed.: Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI 2011), Barcelona, Spain, Morgan Kaufmann (July 2011) 813–818
21. Giordano, L., Gliozzi, V., Olivetti, N., Pozzato, G.L.: A NonMonotonic Description Logic for Reasoning About Typicality. Artificial Intelligence **195** (2013) 165–202
22. Kraus, S., Lehmann, D., Magidor, M.: Nonmonotonic reasoning, preferential models and cumulative logics. Artificial Intelligence **44**(1-2) (1990) 167–207
23. Reiter, R.: A theory of diagnosis from first principles. Artif. Intelligence **32 (1)** (1987) 57–96
24. Giordano, L., Gliozzi, V., Olivetti, N., Pozzato, G.L.: Rational closure in $\mathcal{SHIQ}$. In: DL 2014, 27th International Workshop on Description Logics. Volume 1193 of CEUR Workshop Proceedings., CEUR-WS.org (2014) 543–555
25. Giordano, L., Gliozzi, V.: Encoding a preferential extension of the description logic *SROIQ* into *SROIQ*. In Esposito, F., Pivert, O., Hacid, M., Ras, Z.W., Ferilli, S., eds.: Proc. of Foundations of Intelligent Systems - 22nd Int. Symposium, ISMIS 2015, Lyon, France, October 21-23. Volume 9384 of Lecture Notes in Computer Science., Springer (2015) 248–258