# Rational Defeasible Reasoning for Description Logics


**Katarina Britz** ABRITZ@SUN.AC.ZA
*Centre for Artificial Intelligence Research*
*Stellenbosch University, South Africa*

**Giovanni Casini** GIOVANNI.CASINI@GMAIL.COM
*ICR Research Group*
*Université du Luxembourg*
*Luxembourg*

**Thomas Meyer** TMEYER@CS.UCT.AC.ZA
*Centre for Artificial Intelligence Research*
*CSIR Meraka Institute and University of Cape Town*
*Cape Town, South Africa*

**Kody Moodley** KMOODLEY@CSIR.CO.ZA
*Centre for Artificial Intelligence Research*
*CSIR Meraka Institute and University of KwaZulu-Natal*
*Pretoria, South Africa*

**Uli Sattler** SATTLER@CS.MAN.AC.UK
*Information Management Group*
*University of Manchester*
*Manchester, United Kingdom*

**Ivan Varzinczak** VARZINCZAK@CRIL.FR
*Centre de Recherche en Informatique de Lens*
*Université d'Artois*
*Lens, France*



## Abstract

In this paper, we extend description logics (DLs) with non-monotonic reasoning features. We start by investigating a notion of *defeasible subsumption* in the spirit of defeasible conditionals as studied by Kraus and colleagues in the propositional case. In particular, we consider a natural and intuitive semantics for defeasible subsumption, and we investigate syntactic properties (*à la* Gentzen) for both *preferential* and *rational* subsumptions and prove representation results for the description logic $\mathcal{ALC}$. Such representation results pave the way for more effective decision procedures for defeasible reasoning in DLs. We analyse the problem of non-monotonic reasoning in DL at the level of *entailment* for both TBox and ABox reasoning, and present an adaptation of rational closure for the DL environment. Importantly, we also show that computing it can be reduced to classical $\mathcal{ALC}$ entailment. One of the stumbling blocks to evaluating performance scalability of rational closure is the absence of naturally occurring DL-based ontologies with defeasible features. We overcome this barrier by devising an approach to introduce defeasible subsumption into classical real-world ontologies. Such semi-natural defeasible ontologies, together with a purely artificial set, are used to test our rational closure algorithms. We found that performance is scalable on the whole with no major bottlenecks.

## 1. Introduction

Description logics (DLs) (Baader, Calvanese, McGuinness, Nardi, & Patel-Schneider, 2007) are central to many modern AI and database applications because they provide the logical foundation of formal ontologies. Endowing DLs with defeasible reasoning features is therefore a promising endeavour, drawing on a well-established body of research on non-monotonic reasoning in the field of knowledge representation and reasoning (KR). Indeed, the past 20 years have witnessed many attempts to introduce non-monotonic reasoning capabilities in a DL setting, ranging from preferential approaches (Britz, Heidema, & Meyer, 2008, 2009; Britz, Meyer, & Varzinczak, 2011b; Casini & Straccia, 2010; Giordano, Gliozzi, Olivetti, & Pozzato, 2007, 2008, 2013b, 2015; Quantz & Royer, 1992; Quantz & Ryan, 1993) to circumscription (Bonatti, Lutz, & Wolter, 2009; Bonatti, Faella, & Sauro, 2011a, 2011b; Sengupta, Alfa Krisnadhi, & Hitzler, 2011), amongst others (Baader & Hollunder, 1993, 1995; Donini, Nardi, & Rosati, 2002; Padgham & Zhang, 1993; Bonatti, Faella, Petrova, & Sauro, 2015a; Bonatti, Petrova, & Sauro, 2015b).

Preferential extensions of DLs turn out to be particularly promising, mostly because they are based on an elegant comprehensive and well-studied framework for non-monotonic reasoning in the propositional case proposed by Kraus, Lehmann and Magidor (1990, 1992) and often referred to as the *KLM approach*. Such a framework is valuable for a number of reasons. First, it provides for a thorough analysis of some formal properties that any consequence relation deemed as appropriate in a non-monotonic setting ought to satisfy. Such formal properties, which resemble those of a Gentzen-style proof system, play a central role in assessing how intuitive the obtained results are and enables a more comprehensive characterisation of the introduced non-monotonic conditional from a logical point of view. Second, the KLM approach allows for many decision problems to be reduced to classical entailment checking, sometimes without blowing up the computational complexity with respect to the underlying classical case. Finally, it has a well-known connection with the AGM-approach to belief revision (Gärdenfors & Makinson, 1994; Rott, 2001). It is therefore reasonable to expect that most, if not all, of the aforementioned features of the KLM approach should transfer to KLM-based extensions of DLs.

Following the motivation laid out above, several extensions to the KLM approach to description logics have been proposed recently (Britz et al., 2008, 2011b; Casini & Straccia, 2010; Giordano, Gliozzi, Olivetti, & Pozzato, 2009b; Giordano et al., 2013b, 2015), each of them investigating particular constructions. However, here we shall aim at providing a *comprehensive* study of the formal basis of preferential defeasible reasoning in DLs. By that we mean (*i*) defining a general and intuitive *semantics*; (*ii*) showing a corresponding *representation result* (in the KLM sense of the term) linking the semantics with the KLM-style properties; (*iii*) presenting an appropriate analysis of *entailment* in the context of both TBoxes and ABoxes with defeasible information, and (*iv*) providing an implementation of the underlying theory and an accompanying set of experiments run on sizeable ontologies with defeasible information. To our knowledge, it turns out that none of the existing approaches in the literature accounts for *all* these requirements. Therefore, filling this gap is the main purpose of the present paper.

In the remainder of the paper, we shall take the following route: After fixing the notation (Section 2), we present the notion of defeasible subsumption *à la* KLM (Section 3). In particular, using an intuitive semantics for the idea that "usually, an element of the class $C$ is also an element of the class $D$", we provide a characterisation (via representation results) of two important classes of defeasible statements, namely preferential and rational subsumption. In Section 4, we investigate two obvious candidates for the notion of entailment in the context of defeasible DLs, namely preferential and rank entailment. These turn out not to have all properties seen as important in a non-monotonic DL setting, somehow mimicking a similar feature in the propositional case (Lehmann & Magidor, 1992). As a result we investigate a notion of minimal rank entailment in Section 5. We take this definition further by exploring the relationship that minimal rank entailment has with both Lehmann and Magidor's (1992) definition of rational closure and the more recent algorithm by Casini and Straccia (2010) for its computation (Section 6). In Section 7, we extend this approach to ABox reasoning — importantly, ABoxes enriched with defeasible statements about concept membership. Following that, we present experimental results (Section 8) supporting our claim that preferential description logics are viable in practice. Finally, after a discussion of, and comparison with, related work (Section 9), we conclude with a summary of our contributions and some directions for further exploration. Proofs of our results can be found in the appendix.

## 2. Logical Preliminaries

Description Logics (DLs) (Baader et al., 2007) are a family of logic-based knowledge representation formalisms with interesting computational properties and a variety of applications. In particular, DLs are well-suited for representing and reasoning about terminological knowledge, and constitute the formal foundations of semantic web ontologies. Technically, DLs correspond to decidable fragments of first-order logic and are closely related to modal logics (Schild, 1991). There are many different kinds of description logics with specific expressivity and applications. In this paper, we shall focus on the description logic $\mathcal{ALC}$, which is a main representative of the DL family: it is directly related to normal multi-modal logic $\mathsf{K}$ and it is considered a turning point in DLs from the point of view of expressivity, in the sense of being considered a prototypical compromise between highly expressive DLs and low complexity ones.

The (concept) language of $\mathcal{ALC}$ is built upon a finite set of atomic *concept names* $\mathsf{N}_{\mathscr{C}}$, a finite set of *role names* $\mathsf{N}_{\mathscr{R}}$ and a finite set of *individual names* $\mathsf{N}_{\mathscr{I}}$ such that $\mathsf{N}_{\mathscr{C}}$, $\mathsf{N}_{\mathscr{R}}$ and $\mathsf{N}_{\mathscr{I}}$ are pairwise disjoint. We shall use $A, B, \ldots$, possibly decorated with primes, as 'meta-variables' for the atomic concepts, $r, s, \ldots$, possibly decorated with primes, to denote role names, and $a, b, \ldots$, possibly decorated with primes, to denote individual names. With $C, D, \ldots$, also possibly decorated with primes, we shall denote the complex concepts of our language. These are built using the constructors $\sqcap$ (concept conjunction), $\neg$ (complement), and $\exists$ (existential restriction) according to the following rule

$$C ::= A \mid \neg C \mid (C \sqcap C) \mid \exists r.C$$

Concepts built with the constructors $\sqcup$ and $\forall$ are defined in terms of the others in the usual way. We use $\top$ as an abbreviation for $A \sqcup \neg A$ and $\bot$ as an abbreviation for $A \sqcap \neg A$, for some $A \in \mathsf{N}_{\mathscr{C}}$. With $\mathcal{L}$ we denote the *language* of all $\mathcal{ALC}$ concepts, which is understood as the smallest set of symbol sequences generated according to the rules above. When writing down concepts of $\mathcal{L}$, we shall omit parentheses whenever they are not essential for disambiguation.

The semantics of $\mathcal{ALC}$ is the standard set-theoretic Tarskian semantics. An *interpretation* is a structure $\mathcal{I} := \langle \Delta^{\mathcal{I}}, \cdot^{\mathcal{I}} \rangle$, where $\Delta^{\mathcal{I}}$ is a non-empty set called the *domain*, and $\cdot^{\mathcal{I}}$ is an *interpretation function* mapping concept names $A$ to subsets $A^{\mathcal{I}}$ of $\Delta^{\mathcal{I}}$, role names $r$ to binary relations $r^{\mathcal{I}}$ over $\Delta^{\mathcal{I}}$, and individual names $a$ to elements of the domain $\Delta^{\mathcal{I}}$:

$$A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}},\ r^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}},\ a^{\mathcal{I}} \in \Delta^{\mathcal{I}}$$

Given an interpretation $\mathcal{I} = \langle \Delta^{\mathcal{I}}, \cdot^{\mathcal{I}} \rangle$, $\cdot^{\mathcal{I}}$ is extended to interpret complex concepts in the following way:

$$(\neg C)^{\mathcal{I}} := \Delta^{\mathcal{I}} \setminus C^{\mathcal{I}},\quad (C \sqcap D)^{\mathcal{I}} := C^{\mathcal{I}} \cap D^{\mathcal{I}},$$
$$(\exists r.C)^{\mathcal{I}} := \{x \in \Delta^{\mathcal{I}} \mid \text{for some } y, (x, y) \in r^{\mathcal{I}} \text{ and } y \in C^{\mathcal{I}}\}$$

In particular, for any interpretation $\mathcal{I}$, $\top^{\mathcal{I}} = \Delta^{\mathcal{I}}$, and $\bot^{\mathcal{I}} = \emptyset$.

As an example, let $\mathsf{N}_{\mathscr{C}} \equiv_{\text{def}} \{A_1, A_2, A_3\}$, $\mathsf{N}_{\mathscr{R}} \equiv_{\text{def}} \{r_1, r_2\}$ and $\mathsf{N}_{\mathscr{I}} \equiv_{\text{def}} \{a_1, a_2, a_3\}$. Figure 1 below depicts the DL interpretation $\mathcal{I}_1 = \langle \Delta^{\mathcal{I}_1}, \cdot^{\mathcal{I}_1} \rangle$, where $\Delta^{\mathcal{I}_1} = \{x_i \mid 1 \leq i \leq 9\}$, $A_1^{\mathcal{I}_1} = \{x_1, x_4, x_6\}$, $A_2^{\mathcal{I}_1} = \{x_3, x_5, x_9\}$, $A_3^{\mathcal{I}_1} = \{x_6, x_7, x_8\}$, $r_1^{\mathcal{I}_1} = \{(x_1, x_6), (x_4, x_8), (x_2, x_5)\}$, $r_2^{\mathcal{I}_1} = \{(x_4, x_4), (x_6, x_4), (x_5, x_8), (x_9, x_3)\}$, $a_1^{\mathcal{I}_1} = x_5$, $a_2^{\mathcal{I}_1} = x_1$, $a_3^{\mathcal{I}_1} = x_2$.

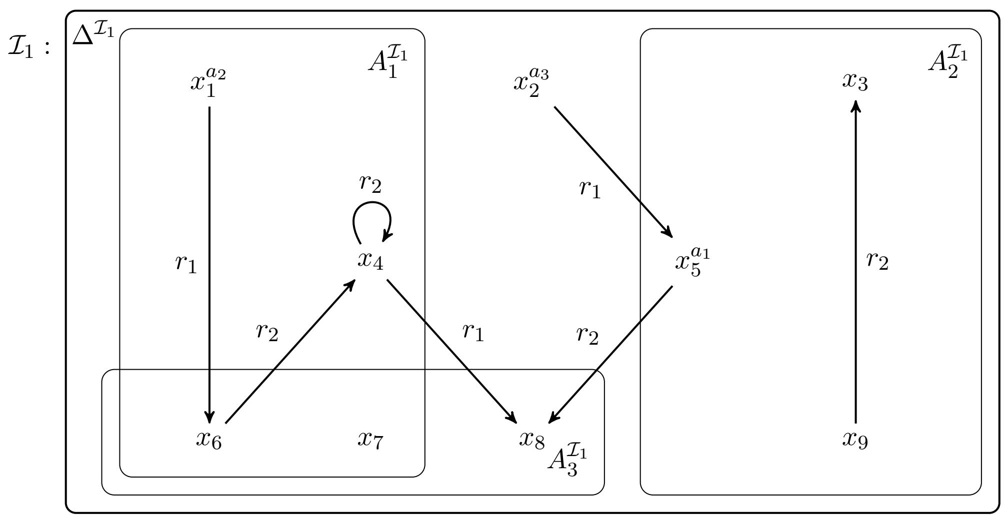


Figure 1: A DL Interpretation for $\mathsf{N}_{\mathscr{C}} = \{A_1, A_2, A_3\}$, $\mathsf{N}_{\mathscr{R}} = \{r_1, r_2\}$ and $\mathsf{N}_{\mathscr{I}} = \{a_1, a_2, a_3\}$.

Given $C, D \in \mathcal{L}$, $C \sqsubseteq D$ is a *subsumption statement*, read "$C$ is subsumed by $D$". $C \equiv D$ is an abbreviation for both $C \sqsubseteq D$ and $D \sqsubseteq C$. An $\mathcal{ALC}$ *TBox* (alias *terminology*) $\mathcal{T}$ is a finite set of subsumption statements. Subsumption statements are also called *general concept inclusion* axioms or GCIs for short.

Given $C \in \mathcal{L}$, $r \in \mathsf{N}_{\mathscr{R}}$ and $a, b \in \mathsf{N}_{\mathscr{I}}$, an *assertional statement* (*assertion*, for short) is an expression of the form $C(a)$ or $r(a, b)$. An $\mathcal{ALC}$ *ABox* (alias *database*) $\mathcal{A}$ is a finite set of assertional statements.

We shall denote statements, both subsumption and assertional, with $\alpha, \beta, \ldots$, possibly with primes.

An interpretation $\mathcal{I}$ *satisfies* a subsumption statement $C \sqsubseteq D$ (denoted $\mathcal{I} \Vdash C \sqsubseteq D$) if $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$. (And then $\mathcal{I} \Vdash C \equiv D$ if $C^{\mathcal{I}} = D^{\mathcal{I}}$.) An interpretation $\mathcal{I}$ *satisfies* an assertion $C(a)$ (respectively, $r(a, b)$), denoted $\mathcal{I} \Vdash C(a)$ (respectively, $\mathcal{I} \Vdash r(a, b)$), if $a^{\mathcal{I}} \in C^{\mathcal{I}}$ (respectively $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in r^{\mathcal{I}}$). Given a statement $\alpha$, with $\models \alpha$ we denote the fact that $\mathcal{I} \Vdash \alpha$ for all interpretations $\mathcal{I}$. We say that $\mathcal{I}$ is a *model* of a TBox $\mathcal{T}$ (denoted $\mathcal{I} \Vdash \mathcal{T}$) if $\mathcal{I} \Vdash \alpha$ for every $\alpha \in \mathcal{T}$. Analogously, $\mathcal{I}$ is a model of an ABox $\mathcal{A}$ (denoted $\mathcal{I} \Vdash \mathcal{A}$) if $\mathcal{I} \Vdash \alpha$ for every $\alpha \in \mathcal{A}$. A statement $\alpha$ is (classically) *entailed* by $\mathcal{T} \cup \mathcal{A}$, denoted $\mathcal{T} \cup \mathcal{A} \models \alpha$, if every model of $\mathcal{T} \cup \mathcal{A}$ satisfies $\alpha$.

For more details on Description Logics in general and on $\mathcal{ALC}$ in particular, the reader is invited to consult the Description Logic Handbook (Baader et al., 2007).

## 3. Preferential Semantics for Description Logics

In this section, we present our semantics for preferential and rational subsumption by enriching standard DL interpretations $\mathcal{I}$ with an ordering on the elements of the domain $\Delta^{\mathcal{I}}$. The intuition underlying it is simple and natural, and extends similar work done for the propositional case by Shoham (1988), Kraus et al. (1990) and Lehmann and Magidor (1992). Variants of the approach we take have been proposed as well by Baltag and Smets (2006, 2008), Boutiller (1994), Britz et al. (2008, 2011b) and Giordano et al. (2007, 2009a, 2009b, 2012, 2013b, 2015). However, as we shall see, this is the first comprehensive semantic account of both preferential and rational subsumption relations based on the standard semantics for description logics.

Informally, our semantic constructions are based on the idea that objects of the domain can be ordered according to their degree of *normality* (Boutilier, 1994) or *typicality* (Booth, Meyer, & Varzinczak, 2012, 2013; Britz et al., 2008; Giordano et al., 2007). Paraphrasing Boutilier (1994, pp. 110–116),

> Surely there is no *inherent* property of objects that allows them to be judged to be more or less normal in absolute terms. These orderings are purely ‘subjective’ (in the sense that they can be thought of as part of an agent’s belief state) and the space of orderings deemed plausible by the agent may (among other things) be determined by e.g. empirical data. By using orderings in this way, we encode our (or the agent’s) *expectations* about the objects corresponding to their perceived regularity or typicality. Those objects not violating our expectations are considered to be more normal than the objects that violate some.

Hence we do not require that there exists something intrinsic about objects that makes one object more normal than another inasmuch as in standard DLs one object need not always be a member of a given concept nor be related with another object via a role. Rather, the intention is to provide a framework in which to express all conceivable ways in which objects, with their associated properties and relationships with other objects, can be ordered in terms of typicality, in the same way that the class of all DL standard interpretations constitute a framework representing all conceivable ways of representing the properties of objects and their relationships with other objects. Just as the latter are constrained by stating subsumption axioms in a knowledge base, the possible orderings that are considered plausible are encoded by what we shall call defeasible subsumption statements (see below).

**Definition 1** [Preferential Interpretation] A preferential interpretation is a structure $\mathcal{P} := \langle \Delta^{\mathcal{P}}, \cdot^{\mathcal{P}}, \prec^{\mathcal{P}} \rangle$, where $\langle \Delta^{\mathcal{P}}, \cdot^{\mathcal{P}} \rangle$ is a DL interpretation (which we denote by $\mathcal{I}_{\mathcal{P}}$ and refer to as the classical interpretation associated with $\mathcal{P}$), and $\prec^{\mathcal{P}}$ is a strict partial order on $\Delta^{\mathcal{P}}$ (i.e., $\prec^{\mathcal{P}}$ is irreflexive and transitive) satisfying the smoothness condition (for every $C \in \mathcal{L}$, if $C^{\mathcal{P}} \neq \emptyset$, then $\min_{\prec^{\mathcal{P}}}(C^{\mathcal{P}}) \neq \emptyset$).[1]

As an example of a preferential interpretation, let $\mathsf{N}_{\mathscr{C}} := \{A_1, A_2, A_3\}$, $\mathsf{N}_{\mathscr{R}} := \{r_1, r_2\}$ and $\mathsf{N}_{\mathscr{I}} := \{a_1, a_2, a_3\}$. Then, Figure 2 below depicts the preferential interpretation $\mathcal{P} = \langle \Delta^{\mathcal{P}}, \cdot^{\mathcal{P}}, \prec^{\mathcal{P}} \rangle$, where $\Delta^{\mathcal{I}} = \{x_i \mid 1 \leq i \leq 8\}$, $A_1^{\mathcal{I}} = \{x_1, x_4, x_6\}$, $A_2^{\mathcal{I}} = \{x_3, x_5, x_8\}$, $A_3^{\mathcal{I}} = \{x_6, x_7\}$, $r_1^{\mathcal{I}} = \{(x_1, x_6), (x_4, x_7), (x_2, x_5)\}$, $r_2^{\mathcal{I}} = \{(x_4, x_4), (x_6, x_4), (x_5, x_7), (x_8, x_3)\}$, $a_1^{\mathcal{I}} = x_5$, $a_2^{\mathcal{I}} = x_1$, $a_3^{\mathcal{I}} = x_8$, and $\prec^{\mathcal{P}}$ is the transitive closure of $\{(x_6, x_4)$, $(x_4, x_1)$, $(x_4, x_2)$, $(x_7, x_4)$, $(x_7, x_5)$, $(x_5, x_2)$, $(x_8, x_5)$, $(x_5, x_3)\}$, i.e., of the relation represented by the dashed arrows in the picture. (Note the direction of the dashed arrows, pointing from more to less preferred objects, with more preferred objects lower in the order.)

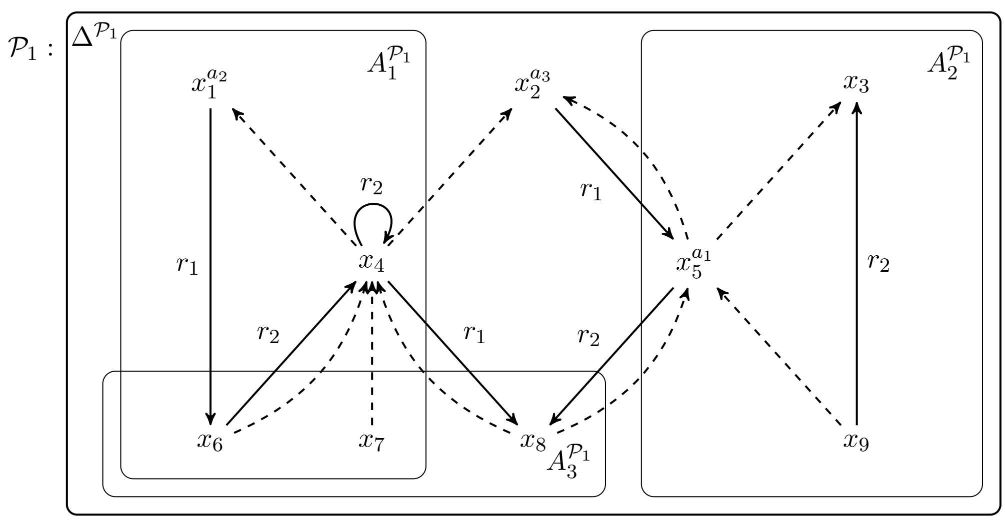


Figure 2: A preferential interpretation.

A preferential interpretation $\mathcal{P}$ *satisfies* a (classical) subsumption statement $C \sqsubseteq D$ (denoted $\mathcal{P} \Vdash C \sqsubseteq D$) if $C^{\mathcal{P}} \subseteq D^{\mathcal{P}}$. It is easy to see that the addition of the $\prec^{\mathcal{P}}$-component preserves the truth of all subsumption statements holding in the remaining structure:

**Observation 1** *Let $\mathcal{P} = \langle \Delta^{\mathcal{P}}, \cdot^{\mathcal{P}}, \prec^{\mathcal{P}} \rangle$ be a preferential interpretation. For every $C, D \in \mathcal{L}$, $\mathcal{P} \Vdash C \sqsubseteq D$ if and only if $\mathcal{I}_{\mathcal{P}} \Vdash C \sqsubseteq D$.*

We are now able to formalise the notion of a *defeasible subsumption*, denoted by $C \mathrel{\sqsubseteq\!\!\!\!\sim} D$.

1. Given $X \subseteq \Delta^{\mathcal{P}}$, with $\min_{\prec^{\mathcal{P}}}(X)$ we denote the set $\{x \in X \mid \text{for every } y \in X, y \not\prec^{\mathcal{P}} x\}$.

**Definition 2** [Satisfaction] Given $C, D \in \mathcal{L}$, a statement of the form $C \sqsubsim D$ is a *defeasible subsumption statement.* A preferential interpretation $\mathcal{P} = \langle \Delta^{\mathcal{P}}, \cdot^{\mathcal{P}}, \prec^{\mathcal{P}} \rangle$ satisfies a defeasible subsumption statement $C \sqsubsim D$, denoted $\mathcal{P} \Vdash C \sqsubsim D$, if $\min_{\prec^{\mathcal{P}}}(C^{\mathcal{P}}) \subseteq D^{\mathcal{P}}$.

A defeasible subsumption statement of the form $C \sqsubsim D$ is to be read as "*usually*, an instance of the concept $C$ is also an instance of the concept $D$". Paraphrasing Lehmann (1989), the intuition of $C \sqsubsim D$ is that "if $C$ were all the information about an object available to an agent, then $D$ would be a sensible conclusion to draw about such an object". It is worth noting that $\sqsubsim$, just as $\sqsubseteq$, is a connective sitting between the concept language (object level) and the meta-language (that of entailment) and it is meant to be the defeasible counterpart of $\sqsubseteq$.

As an example, in the preferential interpretation $\mathcal{P}$ of Figure 2, we have $\mathcal{P} \Vdash A_2 \sqsubsim \exists r_2.\top$ (but note that $\mathcal{P} \not\Vdash A_2 \sqsubseteq \exists r_2.\top$).

**Definition 3** [Defeasible TBox] A *DTBox* (*Defeasible TBox*) $\mathcal{D}$ is a finite set of defeasible subsumption statements $C \sqsubsim D$.

It is worth noting that every (classical) subsumption statement is equivalent, with respect to preferential interpretations, to a defeasible subsumption statement.

**Lemma 1** *For every preferential interpretation $\mathcal{P}$, and every $C, D \in \mathcal{L}$, $\mathcal{P} \Vdash C \sqsubseteq D$ if and only if $\mathcal{P} \Vdash C \sqcap \neg D \sqsubsim \bot$.*

In addition to preferential interpretations we also study *ranked* interpretations, i.e., preferential interpretations in which the $\prec$-component is a *modular* ordering:

**Definition 4** [Modular Order] Given a set $X$, $\prec \subseteq X \times X$ is modular if there is a (strict) totally ordered set $\mathbb{Q}$, with the ordering denoted by $<$, and a ranking function $rk : X \longrightarrow \mathbb{Q}$ s.t. for every $x, y \in X$, $x \prec y$ iff $rk(x) < rk(y)$.

**Definition 5** [Ranked Interpretation] A ranked interpretation is a preferential interpretation $\mathcal{R} = \langle \Delta^{\mathcal{R}}, \cdot^{\mathcal{R}}, \prec^{\mathcal{R}} \rangle$ such that $\prec^{\mathcal{R}}$ is modular.

Intuitively, ranked interpretations allow us to compare any two objects w.r.t. their plausibility. Those with the same rank are technically incomparable, but are viewed as being equally plausible. As such, ranked interpretations are special cases of preferential interpretations, where plausibility can be represented by any strict partial order.

For the remainder of this paper, the reader should keep in mind that it is equivalent to refer to a classical subsumption statement $C \sqsubseteq D$ or to the correspondent *defeasible* subsumption statement $C \sqcap \neg D \sqsubsim \bot$.

In order to study the properties of defeasible subsumption, it is useful to consider certain classes of *defeasible conditionals* of the form $\leadsto \subseteq \mathcal{L} \times \mathcal{L}$ over the concept language $\mathcal{L}$. For clarity of presentation, we make a clear distinction between binary relations over $\mathcal{L}$ satisfying certain structural properties (as below), and classes of defeasible subsumption statements introduced as semantic constructs.

We say that $\leadsto$ is *a preferential conditional* if it satisfies the following set of properties, which we refer to as (the DL versions of the) *preferential* KLM properties:

$$\text{(Cons) } \top \not\leadsto \bot \quad \text{(Ref) } C \leadsto C \quad \text{(LLE) } \frac{\models C \equiv D,\ C \leadsto E}{D \leadsto E}$$

$$\text{(And) } \frac{C \leadsto D,\ C \leadsto E}{C \leadsto D \sqcap E} \qquad \text{(Or) } \frac{C \leadsto E,\ D \leadsto E}{C \sqcup D \leadsto E}$$

$$\text{(RW) } \frac{C \leadsto D,\ \models D \sqsubseteq E}{C \leadsto E} \quad \text{(CM) } \frac{C \leadsto D,\ C \leadsto E}{C \sqcap D \leadsto E}$$

The last six properties are the obvious translations of the properties for preferential consequence relations proposed by Kraus et al. (1990) in the propositional setting. They have been discussed at length in the literature for both the propositional and the DL cases (Kraus et al., 1990; Lehmann & Magidor, 1992; Lehmann, 1995; Giordano et al., 2009b) and we shall not do so here. Property (Cons) corresponds to the requirement that preferential interpretations, like standard DL interpretations, have non-empty domains.

**Definition 6** Let $\mathcal{P}$ be a preferential interpretation, then $\leadsto_{\mathcal{P}} := \{(C, D) \mid \mathcal{P} \Vdash C \precsim D\}$ is the defeasible conditional induced by $\mathcal{P}$.

The first important result we present shows that there is a full correspondence between the class of preferential conditionals and the class of defeasible conditionals induced by preferential interpretations. It is the DL analogue of a result proved by Lehmann et al. for the propositional case (Theorem 3, (Kraus et al., 1990)).

**Theorem 1** *A defeasible conditional $\leadsto\ \subseteq \mathcal{L} \times \mathcal{L}$ is preferential if and only if there is a preferential interpretation $\mathcal{P}$ such that $\leadsto_{\mathcal{P}} = \leadsto$.*[2]

What is perhaps surprising about this result is that no additional properties based on the structure of $\mathcal{ALC}$ are necessary to characterise the conditionals induced by preferential interpretations. We provide below a number of properties involving the use of quantifiers that are satisfied by all preferential conditionals.

The first two are 'existential' and 'universal' versions of cautious monotonicity (CM):

$$(\text{CM}_{\exists}) \ \frac{\exists r.C \leadsto E,\ \exists r.C \leadsto \forall r.D}{\exists r.(C \sqcap D) \leadsto E}$$

$$(\text{CM}_{\forall}) \ \frac{\forall r.C \leadsto E,\ \forall r.C \leadsto \forall r.D}{\forall r.(C \sqcap D) \leadsto E}$$

The third one is a rephrasing of the Rule of Necessitation in modal logic (Chellas, 1980). It guarantees the absence of so-called *spurious objects* (Britz, Meyer, & Varzinczak, 2012) in our original preferential semantics for DLs (Britz, Meyer, & Varzinczak, 2011a; Britz et al., 2011b). That is, if $C$ is inconsistent, then so is $\exists r.C$.

$$\text{(Norm) } \frac{C \leadsto \bot}{\exists r.C \leadsto \bot}$$

2. We remind the reader that all the proofs are given in the Appendix.

If, in addition to the preferential properties, the relation $\rightsquigarrow$ also satisfies rational monotonicity (RM) below, it is said to be a *rational* conditional:

$$\text{(RM)}\ \frac{C \rightsquigarrow E,\ C \not\rightsquigarrow \neg D}{C \sqcap D \rightsquigarrow E}$$

(RM) is considered a desirable property since it is a necessary condition for the satisfaction of *Presumption of Typicality* (Lehmann, 1995, Section 3.1), that is, we reason assuming that we are in the most typical possible situation, compatibly with the information at our disposal (see Section 5). Analogous to the case for cautious monotonicity, the following 'existential' and 'universal' versions of rational monotonicity are satisfied by all rational conditionals:

$$\text{(RM}_\exists\text{)}\ \frac{\exists r.C \rightsquigarrow E,\ \exists r.C \not\rightsquigarrow \forall r.\neg D}{\exists r.(C \sqcap D) \rightsquigarrow E}$$

$$\text{(RM}_\forall\text{)}\ \frac{\forall r.C \rightsquigarrow E,\ \forall r.C \not\rightsquigarrow \forall r.\neg D}{\forall r.(C \sqcap D) \rightsquigarrow E}$$

When considering rational conditionals, one has to move to *ranked* interpretations (Definition 5). This brings us to our second important result, showing that the defeasible conditionals induced by ranked interpretations are precisely the rational conditionals. Again, this is the DL analogue of a result proved by Lehmann and Magidor for the propositional case (Theorem 5, (Lehmann & Magidor, 1992)).

**Theorem 2** *A defeasible conditional* $\rightsquigarrow\ \subseteq \mathcal{L} \times \mathcal{L}$ *is rational if and only if there is a ranked interpretation* $\mathcal{R}$ *such that* $\rightsquigarrow_{\mathcal{R}}\ =\ \rightsquigarrow$.

It is worth pausing for a moment to emphasise the significance of these two results. They provide exact semantic characterisations of two important classes of defeasible conditionals, preferential and rational conditionals, in terms of the classes of preferential and ranked interpretations respectively. As we shall see in section 4, these results form the core of the investigation into appropriate forms of entailment for defeasible knowledge bases.

## 4. Towards Reasoning with Defeasible TBoxes

Given a knowledge base composed of a set of classical subsumptions $C \sqsubseteq D$ and a set of defeasible subsumption statements of the form $C \mathrel{\underset{\sim}{\sqsubset}} D$, from a knowledge representation and reasoning perspective, it becomes important to address the question of what it means for a defeasible subsumption statement to be *entailed* by others.

**Definition 7** [Knowledge Base] A knowledge base is a tuple $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$, where $\mathcal{T}$ is a classical TBox and $\mathcal{D}$ is a defeasible TBox (cf. Definition 3).

Although we will usually be interested in *finite* KBs, in the present section and the next, it is useful to provide a definition for the more general case involving infinite TBoxes and DTBoxes. In a slight abuse of notation, we do not explicitly distinguish between the tuple $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$ and the union $\mathcal{K} = \mathcal{T} \cup \mathcal{D}$, leaving it to the context to clarify the meaning. We use $\alpha, \beta, \ldots$ to denote (classical and defeasible) subsumption statements.

Given a preferential interpretation $\mathcal{P}$ and a KB $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$, we extend the notion of satisfaction to KBs in the obvious way: $\mathcal{P}$ satisfies a knowledge base $\mathcal{K}$ ($\mathcal{P}$ is a model of $\mathcal{K}$), noted as $\mathcal{P} \Vdash \mathcal{K}$, if $\mathcal{P} \Vdash \alpha$ for every statement $\alpha$ in $\mathcal{K}$. $\mathcal{K}$ is said to be *preferentially satisfiable* if there is a preferential interpretation that satisfies $\mathcal{K}$.

The starting point for the definition of an appropriate entailment relation is to consider the Tarskian notion of entailment defined on the basis of *all* the preferential models of a KB.

**Definition 8** [Preferential Entailment] A knowledge base $\mathcal{K}$ preferentially entails a defeasible subsumption statement $\alpha$, denoted $\mathcal{K} \models_P \alpha$, if $\mathcal{P} \Vdash \alpha$ for every $\mathcal{P}$ such that $\mathcal{P} \Vdash \mathcal{K}$.

One of the ways to evaluate versions of defeasible entailment is to consider the defeasible conditional it induces. Note that Definition 9 below applies to a *generic* entailment relation $\models_{\mathcal{X}}$ and not just to $\models_P$.

**Definition 9** The conditional induced by a knowledge base $\mathcal{K}$ under entailment relation $\models_{\mathcal{X}}$ is the set $\leadsto_{\mathcal{X}}^{\mathcal{K}} := \{(C, D) \mid \mathcal{K} \models_{\mathcal{X}} C \mathrel{\sqsubseteq\!\!\!\!\!\sim} D\}$.

The results below are similar to results obtained by Lehmann and Magidor (1992) in the propositional case and support the claim that, within the context of preferential interpretations, preferential entailment is *the unique* appropriate version of entailment.

**Proposition 1** *If a knowledge base $\mathcal{K}$ is preferentially satisfiable, then $\leadsto_P^{\mathcal{K}}$ is preferential.*

Further evidence in favour of preferential entailment is obtained by linking it up to the notion of preferential closure. Given a KB $\mathcal{K}$, we define the *preferential closure* of $\mathcal{K}$ as the intersection of all the preferential conditionals containing the set $\{(C, D) \mid C \mathrel{\sqsubseteq\!\!\!\!\!\sim} D \in \mathcal{D}\} \cup \{(C \sqcap \neg D, \bot) \mid C \sqsubseteq D \in \mathcal{T}\}$.[3]

**Proposition 2** *Let $\leadsto_P^{\mathcal{K}}$ be the (preferential) conditional induced by a KB $\mathcal{K}$ under $\models_P$. Then $\leadsto_P^{\mathcal{K}}$ coincides with the preferential closure of $\mathcal{K}$.*

Given that Definition 8 is Tarskian in nature, it is worth pointing out that $\models_P$ is a *Tarskian consequence relation*, satisfying the following three properties (where $Cn_{\mathcal{X}}(\mathcal{K}) := \{\alpha \mid \mathcal{K} \models_{\mathcal{X}} \alpha\}$, and with $\mathcal{X} = P$):

**(Inclusion)** $\mathcal{K} \subseteq Cn_{\mathcal{X}}(\mathcal{K})$

**(Idempotency)** $Cn_{\mathcal{X}}(\mathcal{K}) = Cn_{\mathcal{X}}(Cn_{\mathcal{X}}(\mathcal{K}))$

**(Monotonicity)** If $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$, $\mathcal{K}' = \langle \mathcal{T}', \mathcal{D}' \rangle$, $\mathcal{T} \subseteq \mathcal{T}'$, $\mathcal{D} \subseteq \mathcal{D}'$, then $Cn_{\mathcal{X}}(\mathcal{K}) \subseteq Cn_{\mathcal{X}}(\mathcal{K}')$

While Inclusion and Idempotency are desirable properties, the Monotonicity property shows that in spite of the non-monotonic features of $\mathrel{\sqsubseteq\!\!\!\!\!\sim}$, we end up with a logic that is *monotonic* at the level of entailment. It can be argued (and has been done in the propositional case (Lehmann & Magidor, 1992; Lehmann, 1995)) that such a property is not an

3. Since the preferential conditionals correspond to the information expressed in defeasible form (i.e., by means of defeasible subsumption statements), we make use of the correspondence between each axiom $C \sqsubseteq D$ in $\mathcal{T}$ and the defeasible axiom $C \sqcap \neg D \mathrel{\sqsubseteq\!\!\!\!\!\sim} \bot$ via Lemma 1.

indication that preferential entailment is an inappropriate version of entailment *within the context of preferential interpretations*, but rather that the class of preferential interpretations is inferentially too weak. It is well-known, for example, that preferential entailment does not support the inheritance of defeasible properties, even in the absence of any conflict (Lehmann, 1995), and such a problem carries over to the case for DLs.

As an example, if we know that both plant cells and mammalian red blood cells are eukaryotic cells ($\mathsf{PlantCell} \sqsubseteq \mathsf{EukCell}$, $\mathsf{MamRBC} \sqsubseteq \mathsf{EukCell}$), that eukaryotic cells usually have a nucleus ($\mathsf{EukCell} \precsim \exists \mathsf{hasNuc}.\top$) and that mammalian red blood cells do not ($\mathsf{MamRBC} \sqsubseteq \neg\exists \mathsf{hasNuc}.\top$), preferential entailment does not allow us to conclude even that plant cells usually have a nucleus ($\mathsf{PlantCell} \precsim \exists \mathsf{hasNuc}.\top$).[4] Hence even the subconcepts that, according to the information at our disposal, behave in a typical way (the plant cells) are not allowed to inherit the typical properties associated to their superclasses (the existence of a nucleus in eukaryotic cells). Such limitations in the inheritance of properties are directly connected to the monotonicity of the entailment relations: assume a monotonic entailment relation (like $\models_P$) and a concept inheriting a typical property of a superconcept, e.g. assume we derive $\mathsf{PlantCell} \precsim \exists \mathsf{hasNuc}.\top$; if we add to the knowledge base the information that plant cells do not have nuclei ($\mathsf{PlantCell} \sqsubseteq \neg\exists \mathsf{hasNuc}.\top$), and hence that plant cells constitute an exceptional subconcept of eukaryotic cells, we would end up with a preferentially consistent knowledge base, but in the meantime we would be forced by monotonicity to conclude also from this knowledge base that $\mathsf{PlantCell} \precsim \exists \mathsf{hasNuc}.\top$, and consequently that plant cells do not exist, since they both should and should not have a nucleus.

It is with this limit in mind that we now shift our attention to defeasible entailment based on *ranked* interpretations. Observe that the notion of satisfiability carries over to ranked interpretations. That is, a knowledge base $\mathcal{K}$ is said to be *rank satisfiable* if there is a ranked interpretation that satisfied $\mathcal{K}$. In fact, it can be shown that preferential and rank satisfaction coincide.

**Proposition 3** *A knowledge base $\mathcal{K}$ is preferentially satisfiable if and only if it is rank satisfiable.*

The first obvious attempt to define defeasible entailment for ranked interpretations is to apply Definition 8 to ranked interpretations.

**Definition 10** [Rank Entailment] A defeasible subsumption statement $\alpha$ is rank entailed by a knowledge base $\mathcal{K}$ (written $\mathcal{K} \models_R \alpha$) if $\mathcal{R} \Vdash \alpha$ for every $\mathcal{R}$ such that $\mathcal{R} \Vdash \mathcal{K}$.

It turns out that rank entailment is still problematic. Firstly, it is clearly a monotonic notion of entailment, just as every notion of entailment which takes under consideration what is satisfied by *all* the models of the premises. Actually, it turns out that it corresponds *exactly* to preferential entailment, as the following result, adapted from a similar result in the propositional case (Lehmann & Magidor, 1992), shows.

**Theorem 3** *A subsumption statement $\alpha$ is preferrentially entailed by a knowledge base $\mathcal{K}$ if and only if it is rank entailed by $\mathcal{K}$. That is,*

$$\mathcal{K} \models_P \alpha \textit{ iff } \mathcal{K} \models_R \alpha.$$

4. This example was originally provided by Piero Bonatti (personal communication).

The result above is easily extendable to suitable definitions of equivalence.

**Definition 11** Two knowledge bases are *preferentially equivalent* (resp., *rank equivalent*) if they have exactly the same preferential models (resp., ranked models).

It follows immediately from Theorem 3 that preferential and rank equivalence coincide.

So rank entailment suffers from exactly the same shortcomings as preferential entailment, including the lack of support for the inheritance of defeasible properties. Related to this is the issue that rank entailment builds on ranked interpretations, but generates a defeasible conditional that is only preferential and does not satisfy RM. In a sense, a commitment to ranked interpretations implies a commitment to rational conditionals (given Theorem 2, which says that the ranked interpretations generate precisely the rational conditionals), but rank entailment violates this commitment.

Observe that, since every classical subsumption statement can be viewed as an abbreviation of a defeasible subsumption statement (Lemma 1), it follows on one hand that a classical TBox can be viewed as a special case of a DTBox (and a knowledge base can be simply expressed using a DTBox). On the other hand, if we want to distinguish between classical information in the TBox and defeasible information in the DTBox, it is possible to have classical information 'hidden' in the DTBox, as the following example illustrates.

**Example 1** Consider a knowledge base $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$, with $\mathcal{T} = \emptyset$ and $\mathcal{D} = \{C \mathrel{\sqsubset\!\!\!\sim} D, C \mathrel{\sqsubset\!\!\!\sim} \neg D\}$. It is easy to see that such a knowledge base is preferentially (and rank) equivalent to the knowledge base $\mathcal{K}' = \langle \mathcal{T}', \mathcal{D}' \rangle$, with $\mathcal{T}' = \{C \sqsubseteq \bot\}$ and $\mathcal{D}' = \emptyset$: Given the validity of the property *And*, we can derive from $\mathcal{D}$ the statement $C \mathrel{\sqsubset\!\!\!\sim} \bot$, which is (preferentially and rank) equivalent to $C \sqsubseteq \bot$; on the other hand, from $C \sqsubseteq \bot$ we have $C \mathrel{\sqsubset\!\!\!\sim} \bot$ and, from the validity of the property RW, we can conclude $C \mathrel{\sqsubset\!\!\!\sim} D$ and $C \mathrel{\sqsubset\!\!\!\sim} \neg D$. Hence the information contained in the DTBox $\mathcal{D}$ is actually classical information 'disguised' as defeasible information.

With the deficiencies of rank entailment in mind, the goal of the rest of this section is to obtain a more appropriate notion of entailment based on ranked interpretations (which we denote by using the generic $\models_{\mathcal{X}}$ symbol). Our purpose here is not to identify a *unique* candidate for entailment for defeasible reasoning, but rather to identify a class of viable candidates. There is a strong argument to be made that, in the case of defeasible reasoning, there is not one unique version of entailment. This argument has been made for the propositional case (Lehmann, 1995), and it carries over to the DL case as well (Casini & Straccia, 2014).

To this end, we start off with some basic desiderata for this generic $\models_{\mathcal{X}}$. Recall that $Cn_{\mathcal{X}}(\mathcal{K}) := \{C \mathrel{\sqsubset\!\!\!\sim} D \mid \mathcal{K} \models_{\mathcal{X}} C \mathrel{\sqsubset\!\!\!\sim} D\}$; our point of departure is then to consider which of the properties of a Tarskian consequence relation, namely Inclusion, Idempotency and Monotonicity, to demand of $Cn_{\mathcal{X}}(\cdot)$.

It seems reasonable to require $\models_{\mathcal{X}}$ to satisfy Inclusion and Idempotency, but, as alluded to above, Monotonicity does not seem appropriate. Recall from our earlier example that, if we know that mammalian red blood cells are eukaryotic ($\mathsf{MamRBC} \sqsubseteq \mathsf{EukCell}$) and eukaryotic cells usually have a nucleus ($\mathsf{EukCell} \mathrel{\sqsubset\!\!\!\sim} \exists \mathsf{hasNuc}.\top$), then we expect that mammalian red blood cells usually have a nucleus ($\mathsf{MamRBC} \mathrel{\sqsubset\!\!\!\sim} \exists \mathsf{hasNuc}.\top$). But on learning that they do not ($\mathsf{MamRBC} \sqsubseteq \neg\exists\mathsf{hasNuc}.\top$), we would expect the conclusion that they usually have a nucleus to be dropped.

The next two requirements relate rank entailment to $\models_{\mathcal{X}}$. The first is based on the idea that, although rank entailment is too weak, it is a suitable lower bound for $\models_{\mathcal{X}}$, representing a *monotonic core* that can be at the base of a more ampliative non-monotonic entailment relation.

**(1)** If $\mathcal{K} \models_R C \sqsubsetsim D$, then $\mathcal{K} \models_{\mathcal{X}} C \sqsubsetsim D$.

The second one is based on the idea that rank entailment deals adequately with *classical* subsumption: reasoning with the classical fragment of the information at our disposal should correspond to using classical entailment, and relying on rank entailment does actually guarantee that.

**(2)** If $\mathcal{K} \models_{\mathcal{X}} C \sqsubseteq D$, then $\mathcal{K} \models_R C \sqsubseteq D$.

Combined with (1) (which implies that, if $\mathcal{K} \models_R C \sqsubsetsim \bot$, then $\mathcal{K} \models_{\mathcal{X}} C \sqsubsetsim \bot$), Requirement (2) insists that the classical consequences of $\models_{\mathcal{X}}$ coincide exactly with the classical consequences of $\models_R$. This is not to say that defeasible subsumption statements play no part in generating classical subsumption statements as consequences (as seen in Example 1, $C \sqsubseteq \bot$ is rank entailed by $\mathcal{D} = \{C \sqsubsetsim D, C \sqsubsetsim \neg D\}$).

The next requirement states the commitment to *rational conditionals* discussed earlier.

**(3)** If $\mathcal{K}$ is satisfiable, then the conditional $\rightsquigarrow$ induced by $\mathcal{K}$ under $\models_{\mathcal{X}}$ is a *rational* conditional.

Requirement (3) dispenses with a number of problems associated with rank entailment, including the inheritance of defeasible properties illustrated by our first example about eukaryotic cells. In fact, as mentioned in Section 3, the satisfaction of (RM) is necessary for reasoning under the *presumption of typicality*, that is immediately related to the inheritance of the defeasible properties of a class by the subclasses not behaving atypically. In the example about eukaryotic cells, we know that eukaryotic cells usually have a nucleus ($\mathsf{EukCell} \sqsubsetsim \exists\mathsf{hasNuc}.\top$) and that plant cells are eukaryotic cells ($\mathsf{PlantCell} \sqsubseteq \mathsf{EukCell}$); rational monotonicity forces us to conclude that, in case we cannot conclude that plant cells are atypical eukaryotic cells ($\mathsf{PlantCell} \not\sqsubsetsim \neg\exists\mathsf{hasNuc}.\top$), we have to treat them as typical eukaryotic cells, that inherit the default properties of eukaryotic cells ($\mathsf{PlantCell} \sqsubsetsim \exists\mathsf{hasNuc}.\top$). In addition, by way of Theorem 2, requirement (3) provides us with a useful technical route to the identification of a suitable $\models_{\mathcal{X}}$. From (3) and Theorem 2 it follows that it is possible to describe $\models_{\mathcal{X}}$ in terms of a *single* ranked interpretation. More precisely, (3) insists that, given any satisfiable knowledge base $\mathcal{K}$, there has to be a ranked interpretation, say $\mathcal{R}^{\mathcal{K}}$, such that for every $C, D \in \mathcal{L}$, $\mathcal{R}^{\mathcal{K}} \Vdash C \sqsubsetsim D$ if and only if $\mathcal{K} \models_{\mathcal{X}} C \sqsubsetsim D$. Our approach will therefore be to identify a suitable $\mathcal{R}^{\mathcal{K}}$ as a *canonical* model for defining entailment whenever $\mathcal{K}$ is satisfiable (if $\mathcal{K}$ is unsatisfiable, we define $\models_{\mathcal{X}}$ such that $\mathcal{K} \models_{\mathcal{X}} C \sqsubsetsim D$ for every $C, D \in \mathcal{L}$). Compliance with (Inclusion) is easy to enforce, we simply require $\mathcal{R}^{\mathcal{K}}$ to satisfy (all elements of) $\mathcal{K}$.

And it is easy to show compliance with (1): if $C \sqsubsetsim D$ is rank entailed by $\mathcal{K}$, then it is satisfied in every ranked interpretation that satisfies $\mathcal{K}$, which includes $\mathcal{R}^{\mathcal{K}}$. As discussed above, it is also easy to see that (3) is satisfied since we know from Theorem 2 that every ranked interpretation generates a rational conditional. Finally, we have also to be sure that the model we choose complies with (2).

In the following section, we will investigate a notion of entailment that satisfies all the above requirements and it is defined by a particular semantic construction.

## 5. Entailment for TBox Reasoning

In this section, we discuss a known instance of entailment for defeasible reasoning that meets all the requirements mentioned above. It is a DL version of the propositional *rational closure* (RC) studied by Lehmann and Magidor (1992). We are going to give in the present section a semantic characterisation; another alternative semantic characterisation of RC in DLs has been proposed by Giordano and others (Giordano, Gliozzi, Olivetti, & Pozzato, 2013a; Giordano et al., 2015), that can be proven to be equivalent to the one presented here. In the following section we will define a procedure that is correct and complete for the present construction and that relies only on classical decision procedures.

Rational closure is a form of inferential closure that is built on the relation of rank entailment $\models_R$, but it extends its inferential power. Such an extension of rank entailment is obtained formalising what is called the *Presumption of Typicality* (Lehmann, 1995, Section 3.1). That is, we always assume that we are dealing with the most typical possible situation, compatible with the information at our disposal. Given the notion of rank entailment, we can build on it a notion of *exceptionality* that is at the base of rational closure.

**Definition 12** [Exceptionality] A concept $C$ is *exceptional* w.r.t. a knowledge base $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$ if $\mathcal{K} \models_R \top \sqsubset\!\!\!\sim \neg C$. An axiom $C \sqsubset\!\!\!\sim D$ is *exceptional* w.r.t. $\mathcal{K}$ if $C$ is exceptional w.r.t. $\mathcal{K}$

So, a concept $C$ is considered exceptional w.r.t. a knowledge base if it is not possible to have a ranked model of the knowledge base in which there is a typical individual (i.e., an individual at least as typical as all the others) that is an instance of the concept $C$. If we apply the notion of exceptionality iteratively, we can associate with every concept $C$ a *ranking value* w.r.t. a knowledge base $\mathcal{K}$ in the following way.

1. We assign a concept $C$ a rank of 0 ($r_{\mathcal{K}}(C) = 0$) if it is not exceptional w.r.t. $\mathcal{K}$, and we set $r_{\mathcal{K}}(C \sqsubset\!\!\!\sim D) = 0$ for every defeasible axiom having $C$ as antecedent. The set of the axioms in $\mathcal{D}$ with rank 0 is denoted as $\mathcal{D}_0^r$.

2. A concept $C$ is assigned a rank of 1 if it does not have a rank of 0 and it is not exceptional w.r.t. the knowledge base $\mathcal{K}^1$ composed by $\mathcal{T}$ and the exceptional part of $\mathcal{D}$, that is, $\mathcal{K}^1 = \langle \mathcal{T}, \mathcal{D} \setminus \mathcal{D}_0^r \rangle$. If $r_{\mathcal{K}}(C) = 1$, then $r_{\mathcal{K}}(C \sqsubset\!\!\!\sim D) = 1$ for every axiom $C \sqsubset\!\!\!\sim D$. The set of the axioms in $\mathcal{D}$ with rank 1 is denoted as $\mathcal{D}_1^r$.

3. In general, for $i > 0$, a concept $C$ is assigned a rank of $i$ if it does not have a rank of $i-1$ and it is not exceptional w.r.t. $\mathcal{K}^i = \langle \mathcal{T}, \mathcal{D} \setminus \bigcup_{j=0}^{i-1} \mathcal{D}_j^r \rangle$. If $r_{\mathcal{K}}(C) = i$, then $r_{\mathcal{K}}(C \sqsubset\!\!\!\sim D) = i$ for every axiom $C \sqsubset\!\!\!\sim D$. The set of the axioms in $\mathcal{D}$ with rank $i$ is denoted $\mathcal{D}_i^r$.

4. By iterating the previous steps, we eventually reach a subset $\mathcal{E} \subseteq \mathcal{D}$ such that all the axioms in $\mathcal{E}$ are exceptional (since $\mathcal{D}$ is finite, we must reach such a point). If $\mathcal{E} \neq \emptyset$, we define the ranking value of the axioms in $\mathcal{E}$ as $\infty$, and the set $\mathcal{E}$ is denoted $\mathcal{D}^r_\infty$.

Following on the procedure above, $\mathcal{D}$ is partitioned into a finite sequence $\langle \mathcal{D}^r_0, \dots, \mathcal{D}^r_n, \mathcal{D}^r_\infty \rangle$ ($n \geq 0$), where $\mathcal{D}^r_\infty$ may be possibly empty. So, through this procedure we can assign a ranking value to every concept and every defeasible subsumption.
Also, using the ranking function we can distinguish between *normally exceptional concepts* and *totally exceptional concepts* w.r.t. a KB: given a defeasible KB $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$, a concept $C$ is *normally exceptional* (resp. *totally exceptional*) w.r.t. $\mathcal{K}$ if $1 \leq r_\mathcal{K} < \infty$ (resp. $r_\mathcal{K} = \infty$). This distinction will be used in Section 8. For a concept $C$ to have $\infty$ as ranking value corresponds to not being satisfiable in any model of $\mathcal{K}$, that is, $\mathcal{K} \models_R C \sqsubseteq \bot$.

Consider the following example, adapted from Giordano et al. (2007):

**Example 2** Let $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$ be a defeasible KB with $\mathcal{T} = \{\mathsf{EmployedStudent} \sqsubseteq \mathsf{Student}\}$ and $\mathcal{D} = \{\mathsf{Student} \mathrel{\sqsubset\!\!\!\sim} \neg\exists \mathsf{receives.TaxInvoice}, \mathsf{EmployedStudent} \mathrel{\sqsubset\!\!\!\sim} \exists \mathsf{receives.TaxInvoice}, \mathsf{EmployedStudent} \sqcap \mathsf{Parent} \mathrel{\sqsubset\!\!\!\sim} \neg\exists \mathsf{receives.TaxInvoice}\}$.

In Example 2, examining the concepts on the LHS of each defeasible subsumption in $\mathcal{K}$, one can verify that $\mathsf{Student}$ is not exceptional w.r.t. $\mathcal{K}$. Therefore, $r_\mathcal{K}(\mathsf{Student}) = 0$. We also find that $r_\mathcal{K}(\mathsf{EmployedStudent}) \neq 0$ and $r_\mathcal{K}(\mathsf{EmployedStudent} \sqcap \mathsf{Parent}) \neq 0$ because both concepts are exceptional w.r.t. $\mathcal{K}$.

$\mathcal{K}^1$ is composed of $\mathcal{T}$ and $\mathcal{D} \setminus \mathcal{D}^r_0$ which, in our example, consists of the defeasible subsumptions in $\mathcal{D}$ except for $\mathsf{Student} \mathrel{\sqsubset\!\!\!\sim} \neg\exists \mathsf{receives.TaxInvoice}$. We find that $\mathsf{EmployedStudent}$ is *not* exceptional w.r.t $\mathcal{K}^1$ and therefore $r_\mathcal{K}(\mathsf{EmployedStudent}) = 1$. Since $\mathsf{EmployedStudent}$ $\sqcap$ $\mathsf{Parent}$ is exceptional w.r.t. $\mathcal{K}^1$, $r_\mathcal{K}(\mathsf{EmployedStudent} \sqcap \mathsf{Parent}) \neq 1$.

Similarly, $\mathcal{K}^2$ is composed of $\mathcal{T}$ and $\{\mathsf{EmployedStudent} \sqcap \mathsf{Parent} \mathrel{\sqsubset\!\!\!\sim} \neg\exists \mathsf{receives.TaxInvoice}\}$. $\mathsf{EmployedStudent} \sqcap \mathsf{Parent}$ is not exceptional w.r.t. $\mathcal{K}^2$ and therefore $r_\mathcal{K}(\mathsf{EmployedStudent} \sqcap \mathsf{Parent}) = 2$.
Adapting Lehmann and Magidor's construction for propositional logic (1992), the rational closure of a knowledge base $\mathcal{K}$ can be defined, referring to ranking values, as follows:

**Definition 13** [Rational Closure] We say that $C \mathrel{\sqsubset\!\!\!\sim} D$ is in the rational closure of a knowledge base $\mathcal{K}$ if

$$r_\mathcal{K}(C \sqcap D) < r_\mathcal{K}(C \sqcap \neg D) \text{ or } r_\mathcal{K}(C) = \infty\ .$$

Informally, the above definition says that $C \mathrel{\sqsubset\!\!\!\sim} D$ is in the rational closure of $\mathcal{K}$ if the ranked models of the knowledge base tell us that some instances of $C \sqcap D$ are more plausible than the instances of $C \sqcap \neg D$.

Applying the definition to the KB in Example 2, we can verify that $\mathsf{Student} \mathrel{\sqsubset\!\!\!\sim} \neg\exists \mathsf{receives.TaxInvoice}$ is in the rational closure of $\mathcal{K}$ because $r_\mathcal{K}(\mathsf{Student} \sqcap \neg\exists \mathsf{receives.TaxInvoice}) = 0$ and $r_\mathcal{K}(\mathsf{Student} \sqcap \exists \mathsf{receives.TaxInvoice}) > 0$. The latter can be derived from the fact that $\mathsf{Student} \sqcap \exists \mathsf{receives.TaxInvoice}$ is exceptional w.r.t. $\mathcal{K}$.

Through analogous arguments, one can derive that both $\mathsf{EmployedStudent} \mathrel{\sqsubset\!\!\!\sim} \exists \mathsf{receives.TaxInvoice}$ and $\mathsf{EmployedStudent} \sqcap \mathsf{Parent} \mathrel{\sqsubset\!\!\!\sim} \neg\exists \mathsf{receives.TaxInvoice}$ are in the rational closure of $\mathcal{K}$ as well.

Returning to the goal of this section, we eventually aim at the definition of an easily implementable decision procedure for rational closure in DLs. Our first step is to describe the construction of a semantic model that is appropriate for the characterisation of rational closure. The semantic construction of rational closure for $\mathcal{ALC}$ that we are going to present is inspired by the semantic characterisation of rational closure for Propositional Logic by Booth and Paris (1998), but in order to adapt it to DLs we had to relevantly depart from the propositional construction. We consider a defeasible knowledge base $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$ and the set $R^{\mathcal{K}}$ of all the ranked models of $\mathcal{K}$. Let $\Delta$ be a countably infinite domain, and $R^{\mathcal{K}}_{\Delta}$ be the set of ranked models of $\mathcal{K}$ that have $\Delta$ as domain. We can prove that the set of models $R^{\mathcal{K}}_{\Delta}$ is sufficient to characterise $\models_R$, i.e., we can rely on a single countably infinite domain $\Delta$ to characterise rank entailment.

**Proposition 4** *For a knowledge base $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$ and a defeasible axiom $C \mathrel{\lesssim} D$, $\mathcal{K} \models_R C \mathrel{\lesssim} D$ if and only if for every $R \in R^{\mathcal{K}}_{\Delta}$, $R \Vdash C \mathrel{\lesssim} D$.*

Note that, since the domain $\Delta$ is countable, for every model in $R^{\mathcal{K}}_{\Delta}$ we can partition the domain $\Delta$ into a sequence of *layers*,

$$\langle L^R_0, \ldots, L^R_n, \ldots \rangle \ ,$$

where for every object $x \in \Delta$, $x \in L^R_0$ iff $x \in \min_{\prec^R}(\Delta)$ and $x \in L^R_{i+1}$ iff $o \in \min_{\prec^R}(\Delta \setminus \bigcup_{0 \leq j \leq i} L^R_j)$. From this partition, we can define the *height* of an individual $a$ as

$$h_R(x) = i \text{ iff } x \in L^R_i \ .$$

The lower the height, the more typical is the individual in the interpretation. We can also think of a level of typicality for the concepts: the *height* of a concept $C$ in an interpretation $R$ ($h_R(C)$) as the lowest (most typical) layer in which the concept's minimal extension is non empty: i.e.

$$h_R(C) = i \text{ iff } \min_{\preceq^R}(C^R) \subseteq L^R_i \ .$$

Now we can use the set of models $R^{\mathcal{K}}_{\Delta}$ to define a model characterising rational closure. Let $R^{\mathcal{K}}_{\cup} = \langle \Delta^{R^{\mathcal{K}}_{\cup}}, \cdot^{R^{\mathcal{K}}_{\cup}}, \prec^{R^{\mathcal{K}}_{\cup}} \rangle$ be a ranked model obtained in the following way.

- For the domain $\Delta^{R^{\mathcal{K}}_{\cup}}$, we consider in $\Delta^{R^{\mathcal{K}}_{\cup}}$ one copy of $\Delta$ for each model in $R^{\mathcal{K}}_{\Delta}$. Specifically, given $\Delta = \{x, y \ldots\}$, we indicate as $\Delta_R = \{x_R, y_R, \ldots\}$ a copy of the domain $\Delta$ associated with an interpretation $R \in R^{\mathcal{K}}_{\Delta}$ and define

$$\Delta^{R^{\mathcal{K}}_{\cup}} = \bigcup_{R \in R^{\mathcal{K}}_{\Delta}} \Delta_R \ .$$

- The interpretation function and the preferential relation are defined referring directly to the models in $R^{\mathcal{K}}_{\Delta}$. That is, for every $x_R, y_{R'} \in \Delta^{R^{\mathcal{K}}_{\cup}}$, every atomic concept $A$ and every role $r$,

  - $x_R \in A^{R^{\mathcal{K}}_{\cup}}$ iff $x \in A^R$;

– $(x_R, y_{R'}) \in r^{R_\cup^\mathcal{K}}$ iff $R = R'$ and $(x, y) \in r^R$;
– $x_R \prec^{R_\cup^\mathcal{K}} y_{R'}$ iff $h_R(x) < h_{R'}(y)$.

It is easy to check by induction on the construction of the concepts that, for every $x_R \in \Delta^{R_\cup^\mathcal{K}}$ and every concept $C$,

– $x_R \in C^{R_\cup^\mathcal{K}}$ iff $x \in C^R$;

and that every individual $x_R \in \Delta^{R_\cup^\mathcal{K}}$ preserves its original height, that is,

– $h_{R_\cup^\mathcal{K}}(x_R) = h_R(x)$.

It turns out that $R_\cup^\mathcal{K}$ is a ranked model characterising the Rational Closure of $\mathcal{K}$.

**Theorem 4** *Let $\mathcal{K}$ be a knowledge base having a ranked model. Then $R_\cup^\mathcal{K}$ is a model of $\mathcal{K}$*[5] *and for any pair of concepts $C, D$, $R_\cup^\mathcal{K} \Vdash C \mathrel{\sqsubset\!\!\!\sim} D$ if and only if $r_\mathcal{K}(C \sqcap D) < r_\mathcal{K}(C \sqcap \neg D)$ or $r(C) = \infty$.*

And an immediate consequence of Theorem 4 and Definition 13 is that for every knowledge base $\mathcal{K}$, $R_\cup^\mathcal{K}$ is a characteristic model of the rational closure of $\mathcal{K}$. From now on, we shall use the symbol $\models_R^\cup$ to indicate this notion of entailment. That is, $\mathcal{K} \models_R^\cup \mathcal{C} \mathrel{\sqsubset\!\!\!\sim} D$ if and only $R_\cup^\mathcal{K} \Vdash C \mathrel{\sqsubset\!\!\!\sim} D$.

**Corollary 1** *$C \mathrel{\sqsubset\!\!\!\sim} D$ is in the rational closure of a knowledge base $\mathcal{K}$ if and only if $\mathcal{K} \models_R^\cup C \mathrel{\sqsubset\!\!\!\sim} D$ (if and only if $R_\cup^\mathcal{K} \Vdash C \mathrel{\sqsubset\!\!\!\sim} D$).*

An alternative semantic characterisation of rational closure in DLs has been presented by Giordano et al. (2013b, 2015). Since their definition of rational closure is identical to ours, it follows that the two semantic characterisations produce identical results.

## 6. Computing Rational Closure

So far, we have shown that rational closure satisfies a number of desirable criteria for rational defeasible entailment. We now present an algorithm to compute rational closure for an $\mathcal{ALC}$ knowledge base. The advantages of our approach, which modifies the algorithm presented by Casini and Straccia (2010) (such an algorithm needed to be slightly modified since it does not always give back the correct result in case $\mathcal{D}_\infty^r \neq \emptyset$), are that it relies completely on classical $\mathcal{ALC}$-entailment, is easily implementable, and has computational complexity that is no worse than that of classical $\mathcal{ALC}$-entailment. We shall use the symbol $\vdash_r$ to indicate the inference relation defined by the algorithm. The aims of this section are therefore to present the decision procedure defining $\vdash_r$, and to prove the correspondence between $\models_R^\cup$ and $\vdash_r$.

Consider a knowledge base $\langle \mathcal{T}, \mathcal{D} \rangle$, with a finite TBox $\mathcal{T} = \{E_1 \sqsubseteq F_1, \ldots, E_m \sqsubseteq F_m\}$ and a finite DTBox $\mathcal{D} = \{C_1 \mathrel{\sqsubset\!\!\!\sim} D_1, \ldots, C_n \mathrel{\sqsubset\!\!\!\sim} D_n\}$. The algorithm to define the inference relation $\vdash_r$ is as follows:

5. The result that $R_\cup^\mathcal{K}$ is a ranked model for $\mathcal{K}$ is based on the closure of $\mathcal{ALC}$ models under disjoint union. This means that our exposition could be extended to other DLs which enjoy this property (e.g. $\mathcal{SHIQ}$).

**Step 1.** The first step assigns a rank to each axiom in the DTBox $\mathcal{D}$.

Central to the algorithm is the exceptionality function $E_{\mathcal{T}}$, aimed at modelling the semantic notion of exceptionality in Definition 12 and formalised by Procedure Exceptional presented below. Given a set of defeasible axioms $\mathcal{E} \subseteq \mathcal{D}$, the procedure returns a subset $\mathcal{E}'$ of $\mathcal{E}$ such that $\mathcal{E}'$ is exceptional w.r.t. $\langle \mathcal{T}, \mathcal{E} \rangle$.

---

**Procedure** Exceptional($\mathcal{T}, \mathcal{E}$)

---

**Input**: $\mathcal{T}$ and $\mathcal{E} \subseteq \mathcal{D}$
**Output**: $\mathcal{E}' \subseteq \mathcal{E}$ such that $\mathcal{E}'$ is exceptional w.r.t. $\mathcal{E}$

```
1 E' := ∅;
2 foreach C ⊏~ D ∈ E do
3     if T ⊨ ⊓ Ē ⊑ ¬C then
4         E' := E' ∪ {C ⊏~ D};
5 return E'
```

---

If we provide the KB in Example 2 as input to Procedure Exceptional, we obtain the output $\mathcal{D}' = \{\mathsf{EmployedStudent} \mathrel{\sqsubset\!\!\!\sim} \exists\mathsf{receives.TaxInvoice}, \mathsf{EmployedStudent} \sqcap \mathsf{Parent} \mathrel{\sqsubset\!\!\!\sim} \neg\exists\mathsf{receives.TaxInvoice}\}$. This is because both concepts on the LHS of the subsumptions in $\mathcal{D}'$ are exceptional (in the semantic sense) w.r.t. the original KB $\mathcal{K}$ in Example 2.

The symbol $\overline{\mathcal{E}}$ in the exceptionality procedure represents the *materialisation* of the set $\mathcal{E}$, that is, the set containing the concepts $\neg C \sqcup D$ that represent at the local level the same inclusion as an axiom $C \mathrel{\sqsubset\!\!\!\sim} D$. That is,

$$\overline{\mathcal{E}} = \{\neg C \sqcup D \mid C \mathrel{\sqsubset\!\!\!\sim} D \in \mathcal{E}\}.$$

This construction helps to define a reduction of concept exceptionality to classical DL entailment (Line 3 of Procedure Exceptional).

That is, a concept $C$ (and an axiom $C \mathrel{\sqsubset\!\!\!\sim} D$) is exceptional w.r.t. $\langle \mathcal{T}, \mathcal{E} \rangle$ if and only if

$$\mathcal{T} \models \bigsqcap \overline{\mathcal{E}} \sqsubseteq \neg C.$$

We shall later prove (Proposition 7 on page 23, proved in the Appendix) that this reduction, as embedded in the presented procedures, captures the notion of exceptionality we defined earlier (Definition 12 on page 14) at the semantic level.

We now define the overall ranking algorithm, presented in Procedure ComputeRanking below. The procedure consists of a finite sequence of applications of the exceptionality procedure, starting from the knowledge base $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$. The algorithm stops once all the classical information possibly 'hidden' inside the DTBox have been moved to the TBox, and we obtain also a complete ranking of the axioms in the DTBox.

**Step 1.a** In brief, we start by setting $\mathcal{T}^* := \mathcal{T}$, $\mathcal{D}^* := \mathcal{D}$ (This corresponds to lines 1-2 of Procedure ComputeRanking). In Example 2, $\mathcal{T}^* := \{\ \mathsf{EmployedStudent} \sqsubseteq \mathsf{Student}\}$ and $\mathcal{D}^* := \{\mathsf{Student} \mathrel{\sqsubset\!\!\!\sim} \neg\exists\mathsf{receives.TaxInvoice}, \mathsf{EmployedStudent} \mathrel{\sqsubset\!\!\!\sim} \exists\mathsf{receives.TaxInvoice},$

**Procedure** ComputeRanking($\mathcal{K}$)

**Input**: KB $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$

**Output**: KB $\langle \mathcal{T}^*, \mathcal{D}^* \rangle$ and the partitioning (ranking) $R = \{\mathcal{D}_0, \ldots, \mathcal{D}_n\}$ for $\mathcal{D}^*$

1 $\mathcal{T}^* := \mathcal{T}$;<br>
2 $\mathcal{D}^* := \mathcal{D}$;<br>
3 $R := \emptyset$;<br>
4 **repeat**<br>
5     $i := 0$;<br>
6     $\mathcal{E}_0 := \mathcal{D}^*$;<br>
7     $\mathcal{E}_1 := \texttt{Exceptional}(\mathcal{T}^*, \mathcal{E}_0)$;<br>
8     **while** $\mathcal{E}_{i+1} \neq \mathcal{E}_i$ **do**<br>
9         $i := i + 1$;<br>
10         $\mathcal{E}_{i+1} := \texttt{Exceptional}(\mathcal{T}^*, \mathcal{E}_i)$;<br>
11     $\mathcal{D}^*_\infty := \mathcal{E}_i$;<br>
12     $\mathcal{T}^* := \mathcal{T}^* \cup \{C \sqsubseteq D \mid C \mathrel{\underset{\sim}{\sqsubset}} D \in \mathcal{D}_\infty\}$;<br>
13     $\mathcal{D}^* := \mathcal{D}^* \setminus \mathcal{D}^*_\infty$;<br>
14 **until** $\mathcal{D}^*_\infty = \emptyset$;<br>
15 **for** $j = 1$ *to* $i$ **do**<br>
16     $\mathcal{D}_{j-1} := \mathcal{E}_{j-1} \setminus \mathcal{E}_j$;<br>
17     $R := R \cup \{\mathcal{D}_{j-1}\}$;<br>
18 **return** $\langle \mathcal{T}^*, \mathcal{D}^* \rangle, R$;

$\mathsf{EmployedStudent} \sqcap \mathsf{Parent} \mathrel{\underset{\sim}{\sqsubset}} \neg\exists \mathsf{receives.TaxInvoice}\}$. We then repeatedly invoke Procedure Exceptional to obtain a sequence of sets of defeasible axioms $\mathcal{E}_0, \mathcal{E}_1, \ldots$, where $\mathcal{E}_0 = \mathcal{D}^*$ and $\mathcal{E}_{i+1}$ is a set of exceptional axioms in $\mathcal{E}_i$ (Lines 4 - 14 of Procedure ComputeRanking).

Let $\overline{\mathcal{D}^*}$ be the materialization of $\mathcal{D}^*$, i.e., $\overline{\mathcal{D}^*} = \{\neg C \sqcup D \mid C \mathrel{\underset{\sim}{\sqsubset}} D \in \mathcal{D}\}$. Further, let $\mathfrak{A}_{\mathcal{D}^*}$ be the set of the antecedents of axioms in $\mathcal{D}^*$, i.e., $\mathfrak{A}_{\mathcal{D}^*} = \{C \mid C \mathrel{\underset{\sim}{\sqsubset}} D \in \mathcal{D}^*\}$). We determine an *exceptionality ranking* of the axioms in $\mathcal{D}^*$ using Procedure Exceptional, $\mathfrak{A}_{\mathcal{D}^*}$, $\mathcal{T}^*$ and $\overline{\mathcal{D}^*}$. That is, we start by checking, for every concept $C$ in $\mathfrak{A}_{\mathcal{D}^*}$, whether

$$\mathcal{T}^* \models \bigsqcap \overline{\mathcal{D}^*} \sqsubseteq \neg C.$$

In case $C$ is exceptional, every axiom $C \mathrel{\underset{\sim}{\sqsubset}} D \in \mathcal{D}^*$ is exceptional w.r.t. $\langle \mathcal{T}^*, \mathcal{D}^* \rangle$ and collectively form the set $\mathcal{E}_1$. If $\mathcal{E}_1 \neq \emptyset$, we repeat the exceptionality procedure for $\langle \mathcal{T}^*, \mathcal{E}_1 \rangle$ defining the set $\mathcal{E}_2$, and so on. We indicate with $E_{\mathcal{T}}$ the function that associates with a given DTBox $\mathcal{D}$ the subset of $\mathcal{D}$ consisting of those axioms that are exceptional w.r.t. $\langle \mathcal{T}, \mathcal{D} \rangle$.

Hence, given $\langle \mathcal{T}^*, \mathcal{D}^* \rangle$, we can construct a sequence $\mathcal{E}_0, \mathcal{E}_1, \ldots$ in the following way:

- $\mathcal{E}_0 := \mathcal{D}^*$
- $\mathcal{E}_{j+1} := E_{\mathcal{T}^*}(\mathcal{E}_j)$

Continuing with Example 2, we obtain the following exceptionality sequence: $\mathcal{E}_0 := \{\mathsf{Student} \precsim \neg\exists\mathsf{receives.TaxInvoice}, \mathsf{EmployedStudent} \precsim \exists\mathsf{receives.TaxInvoice}, \mathsf{EmployedStudent} \sqcap \mathsf{Parent} \precsim \neg\exists\mathsf{receives.TaxInvoice}\}$, $\mathcal{E}_1 := \{\mathsf{EmployedStudent} \precsim \exists\mathsf{receives.TaxInvoice}, \mathsf{EmployedStudent} \sqcap \mathsf{Parent} \precsim \neg\exists\mathsf{receives.TaxInvoice}\}$ and $\mathcal{E}_2 := \{\mathsf{EmployedStudent} \sqcap \mathsf{Parent} \precsim \neg\exists\mathsf{receives.TaxInvoice}\}$.

Since $\mathcal{D}^*$ is finite, the construction will terminate with an empty or non-empty fixed point $\mathcal{E}_i$ of $E_{\mathcal{T}^*}$, that is, a set $\mathcal{E}_i$ such that $E_{\mathcal{T}^*}(\mathcal{E}_i) = \mathcal{E}_i$. In the case of Example 2, we end up with $\mathcal{E}_3 = \emptyset$ as the empty fixed point. However, if this fixed point were to contain axioms, then these axioms are said to have infinite rank (hence we would set $\mathcal{D}^*_\infty := \mathcal{E}_i$ on Line 11 of Procedure ComputeRanking), and the classical translations of these axioms should be moved to the TBox. Hence we redefine the knowledge base in the following way (Lines 12 - 13 of Procedure ComputeRanking):

- $\mathcal{D}^* := \mathcal{D}^* \setminus \mathcal{D}^*_\infty$;
- $\mathcal{T}^* := \mathcal{T}^* \cup \{C \sqsubseteq D \mid C \precsim D \in \mathcal{D}^*_\infty\}$.

Again, for Example 2, we do not have axioms of infinite rank, so this case does not apply. However, in general, we would have to repeat the above procedure (Lines 4 - 14 of Procedure ComputeRanking) until $\mathcal{D}^*_\infty$ is empty (that is, until we have moved all the classical knowledge 'hidden' in the DTBox to the TBox).

Procedure ComputeRanking must terminate since $\mathcal{D}$ is finite, and at every reiteration $\mathcal{D}^*$ becomes smaller (hence, we have at most $|\mathcal{D}|$ reiterations). In the end, we obtain a knowledge base $\langle \mathcal{T}^*, \mathcal{D}^* \rangle$ which is rank equivalent to the original knowledge base $\langle \mathcal{T}, \mathcal{D} \rangle$ (see Proposition 6 below), in which $\mathcal{D}^*$ has no axioms of infinite rank (all the strict knowledge "hidden" in the DTBox has been moved to the TBox). In the following, we say that such a knowledge base is in *rank normal form*.

**Step 1.b** Once we have obtained the knowledge base $\langle \mathcal{T}^*, \mathcal{D}^* \rangle$ and the final sequence $\mathcal{E}_0, \mathcal{E}_1, \ldots, \mathcal{E}_i$, we define a ranking function $r$ that associates to every axiom in $\mathcal{D}^*$ a number, representing its level of exceptionality:

$$r(C \precsim D) = \begin{cases} j & \text{if } C \precsim D \in \mathcal{E}_j \text{ and } C \precsim D \notin \mathcal{E}_{j+1} \ (j < i) \\ \infty & \text{if } C \precsim D \in \mathcal{E}_i \end{cases}$$

And the same applies to the concepts appearing as antecedents in the axioms in $\mathcal{D}$:

$$r(C) = r(C \precsim D) \text{ for every } C \precsim D \in \mathcal{D}^*$$

We indicate with $\mathcal{D}_j$ the set of the defeasible axioms in $\mathcal{D}^*$ having $j$ as ranking value. Hence the set $\mathcal{D}^*$ is partitioned into the sets $\mathcal{D}_0, \ldots, \mathcal{D}_n$, for some $n \geq 0$ (Lines 15 - 17 of Procedure ComputeRanking). For the KB in Example 2, we obtain the sequence: $\mathcal{D}_0 := \{\mathsf{Student} \precsim \neg\exists\mathsf{receives.TaxInvoice}\}$, $\mathcal{D}_1 := \{\mathsf{EmployedStudent} \precsim \exists\mathsf{receives.TaxInvoice}\}$ and $\mathcal{D}_2 := \{\mathsf{EmployedStudent} \sqcap \mathsf{Parent} \precsim \neg\exists\mathsf{receives.TaxInvoice}\}$. This sequence represents the exceptionality ranking of the KB in Example 2.

Now we have moved all the classical information possibly 'hidden' inside the DTBox into the TBox, and we have given a ranking value to all the remaining defeasible axioms.

Note that for every $j, i \leq n$, if $j \leq i$, then $\models \bigsqcap \overline{\mathcal{E}_j} \sqsubseteq \bigsqcap \overline{\mathcal{E}_i}$. Note also that the $\mathcal{D}_j$s are pairwise disjoint: for every $i, j \leq n$, $\mathcal{D}_i \cap \mathcal{D}_j = \emptyset$.

**Step 2.** So, once and for all, we 'prepare' our initial ontology $\langle \mathcal{T}, \mathcal{D} \rangle$, obtaining a rank equivalent ontology $\langle \mathcal{T}^*, \mathcal{D}^* \rangle$ (see Proposition 6 below) and the ranking of the defeasible axioms. Having done that, we can now define the decision procedure, that is, the algorithm to decide whether or not a subsumption statement $C \precsim D$ is in the rational closure of $\langle \mathcal{T}, \mathcal{D} \rangle$, that is, we can define the inference relation $\vdash_r$.

To do that, we use the same approach used in Procedure Exceptional, that is, given $\langle \mathcal{T}^*, \mathcal{D}^* \rangle$ and our sequence of sets $\mathcal{E}_0, \ldots, \mathcal{E}_n$, we use the TBox $\mathcal{T}^*$ and the sets of conjunctions of materialisations $\bigsqcap \overline{\mathcal{E}_0}, \ldots, \bigsqcap \overline{\mathcal{E}_n}$.

We can extend the ranking function $r$ to all the concepts and all the defeasible axioms in our language in the following way. For all concepts $C$ and $D$:

- $r(C) = i$, $0 \leq i \leq n$, iff $\bigsqcap \overline{\mathcal{E}_i}$ is the first element in $\langle \bigsqcap \overline{\mathcal{E}_0}, \ldots, \bigsqcap \overline{\mathcal{E}_n} \rangle$ s.t. $\mathcal{T}^* \not\models \bigsqcap \overline{\mathcal{E}_i} \sqcap C \sqsubseteq \bot$;
- $r(C) = \infty$ iff there is no such $\bigsqcap \overline{\mathcal{E}_i}$;
- $r(C \precsim D) = r(C)$.

It is immediate to see that for every $C$ s.t. $C \precsim D \in \mathcal{D}^*$, this definition of $r$ returns exactly the same result as the $r$ defined in **Step 1.b**.

**Definition 14** $\langle \mathcal{T}, \mathcal{D} \rangle \vdash_r C \precsim D$ if and only if $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_i} \sqcap C \sqsubseteq D$, where $\bigsqcap \overline{\mathcal{E}_i}$ is the first element of the sequence $\bigsqcap \overline{\mathcal{E}_0}, \ldots, \bigsqcap \overline{\mathcal{E}_n}$ such that $\mathcal{T}^* \not\models \bigsqcap \overline{\mathcal{E}_i} \sqsubseteq \neg C$. If there is no such element, $\langle \mathcal{T}, \mathcal{D} \rangle \vdash_r C \precsim D$ if and only if $\mathcal{T}^* \models C \sqsubseteq D$.

Observe that $\langle \mathcal{T}, \mathcal{D} \rangle \vdash_r C \sqsubseteq D$ if and only if $\langle \mathcal{T}, \mathcal{D} \rangle \vdash_r C \sqcap \neg D \precsim \bot$, i.e., if and only if $\mathcal{T}^* \models C \sqcap \neg D \sqsubseteq \bot$ (that is to say, $\mathcal{T}^* \models C \sqsubseteq D$).

The algorithm corresponding to **Step 2** is presented in Procedure RationalClosure. For Example 2, if we would like to check if $\mathsf{EmployedStudent} \precsim \exists \mathsf{receives.TaxInvoice}$ is in the rational closure of $\mathcal{K}$, then we can verify that the while-loop on Line 2 of Procedure RationalClosure terminates when $i = 1$. At this stage, $\bigsqcap \overline{\mathcal{E}_i} = (\neg \mathsf{EmployedStudent} \sqcup \exists \mathsf{receives.TaxInvoice}) \sqcap (\neg \mathsf{EmployedStudent} \sqcup \neg \mathsf{Parent} \sqcup \neg \exists \mathsf{receives.TaxInvoice})$.

One can verify from this that $\mathcal{T}^* \not\models \bigsqcap \overline{\mathcal{E}_i} \sqcap C \sqsubseteq \bot$, i.e., that $\{\ \mathsf{EmployedStudent} \sqsubseteq \mathsf{Student}\ \} \not\models (\neg \mathsf{EmployedStudent} \sqcup \exists \mathsf{receives.TaxInvoice}) \sqcap (\neg \mathsf{EmployedStudent} \sqcup \neg \mathsf{Parent} \sqcup \neg \exists \mathsf{receives.TaxInvoice}) \sqcap \mathsf{EmployedStudent} \sqsubseteq \bot$.

Finally, it is easy to confirm that $\mathcal{T}^* \not\models \bigsqcap \overline{\mathcal{E}_i} \sqcap C \sqsubseteq D$, i.e., that $\{\mathsf{EmployedStudent} \sqsubseteq \mathsf{Student}\ \} \not\models (\neg \mathsf{EmployedStudent} \sqcup \exists \mathsf{receives.TaxInvoice}) \sqcap (\neg \mathsf{EmployedStudent} \sqcup \neg \mathsf{Parent} \sqcup \neg \exists \mathsf{receives.TaxInvoice}) \sqcap \mathsf{EmployedStudent} \sqsubseteq \exists \mathsf{receives.TaxInvoice}$.

Let $C_r$ be the closure operation corresponding to $\vdash_r$, i.e., $C_r(\langle \mathcal{T}, \mathcal{D} \rangle) = \{C \precsim D \mid \langle \mathcal{T}, \mathcal{D} \rangle \vdash_r C \precsim D\} \cup \{C \sqsubseteq D \mid \langle \mathcal{T}, \mathcal{D} \rangle \vdash_r C \sqcap \neg D \precsim \bot\}$. Next, we state that $\mathcal{T}$ and $\mathcal{D}$

**Procedure** RationalClosure($\mathcal{K}$)

**Input**: KB $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$, the correspondent knowledge base $\langle \mathcal{T}^*, \mathcal{D}^* \rangle$, the sequence $\mathcal{E}_0, \ldots, \mathcal{E}_n$, a query $C \tilde{\sqsubseteq} D$.
**Output**: `true` iff $\mathcal{K} \vdash_r C \tilde{\sqsubseteq} D$

$i := 0$;
**while** $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_i} \sqcap C \sqsubseteq \bot$ *and* $i \leq n$ **do**
  $i := i + 1$;
**if** $i \leq n$ **then**
  **return** $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_i} \sqcap C \sqsubseteq D$;
**else**
  **return** $\mathcal{T}^* \models C \sqsubseteq D$;

are in $C_r(\langle \mathcal{T}, \mathcal{D} \rangle)$ and that $C_r(\langle \mathcal{T}, \mathcal{D} \rangle)$ defines a rational conditional (the proof for Proposition 5 is given in Appendix D).

**Proposition 5** *Given a knowledge base $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$, $\mathcal{T} \cup \mathcal{D} \subseteq C_r(\langle \mathcal{T}, \mathcal{D} \rangle)$. Moreover, $C_r(\mathcal{K})$ defines a defeasible conditional $\leadsto_r^{\mathcal{K}}$ that is rational, where $\leadsto_r^{\mathcal{K}} := \{(C, D) \mid \mathcal{K} \vdash_r C \tilde{\sqsubseteq} D\}$.*

The preferential entailment satisfies also two stronger forms of (LLE) and (RW), viz.(LLE') and (RW'):

$$\text{(RW')}\ \frac{\langle \mathcal{T}, \mathcal{D} \rangle \models_{\mathcal{P}} C \tilde{\sqsubseteq} D,\ \mathcal{T} \models D \sqsubseteq E}{\langle \mathcal{T}, \mathcal{D} \rangle \models_{\mathcal{P}} C \tilde{\sqsubseteq} E} \qquad \text{(LLE')}\ \frac{\mathcal{T} \models C \equiv D,\ \langle \mathcal{T}, \mathcal{D} \rangle \models_{\mathcal{P}} C \tilde{\sqsubseteq} E}{\langle \mathcal{T}, \mathcal{D} \rangle \models_{\mathcal{P}} D \tilde{\sqsubseteq} E}$$

The proofs for this are quite straightforward, since all the preferential models satisfying $\langle \mathcal{T}, \mathcal{D} \rangle$ must satisfy $\mathcal{T}$. The validity of the analogous properties w.r.t. $\vdash_r$ can be shown similarly for (RW) and (LLE).

There are some more properties that must be proved in order to verify that $\vdash_r$ corresponds to the notion of rational closure entailment presented in the previous section (Definition 13).

The following lemma states that, as in the propositional case (Lehmann & Magidor, 1992), our procedure correctly manages the classical information, that is, an axiom $C \tilde{\sqsubseteq} \bot$ is in the rational closure of $\langle \mathcal{T}, \mathcal{D} \rangle$ if and only if it is also a consequence of $\langle \mathcal{T}, \mathcal{D} \rangle$ according to rank entailment.

**Lemma 2** *Assume that $C \tilde{\sqsubseteq} D \in \mathcal{D}$. $\langle \mathcal{T}, \mathcal{D} \rangle \models_R C \tilde{\sqsubseteq} \bot$ if and only if $r(C) = \infty$ if and only if $\mathcal{T}^* \models C \sqsubseteq \bot$.*

An immediate consequence of Lemma 2 binds preferential consistency to classical consistency.

**Corollary 2** *$\langle \mathcal{T}, \mathcal{D} \rangle \models_R \top \tilde{\sqsubseteq} \bot$ if and only if $\mathcal{T}^* \models \top \sqsubseteq \bot$.*

We can now prove that the knowledge bases $\langle \mathcal{T}, \mathcal{D} \rangle$ and $\langle \mathcal{T}^*, \mathcal{D}^* \rangle$ (in rank normal form) are rank equivalent.

**Proposition 6** *Consider a knowledge base $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$ and the knowledge base $\mathcal{K}^* = \langle \mathcal{T}^*, \mathcal{D}^* \rangle$ obtained from $\mathcal{K}$ using Procedure ComputeRanking. $\mathcal{K}$ and $\mathcal{K}^*$ are rank equivalent.*

Now we are justified in using the rank normal form $\langle \mathcal{T}^*, \mathcal{D}^* \rangle$ in order to analyse the rational closure of the knowledge base $\langle \mathcal{T}, \mathcal{D} \rangle$. Hence, in what follows we will assume that the knowledge bases $\langle \mathcal{T}, \mathcal{D} \rangle$ we are working with are already in rank normal form (hence they do not have an infinite rank, that is $\mathcal{D}_\infty = \emptyset$). Proposition 6 allows us to make such a step without any harm since we can assume that our knowledge base $\langle \mathcal{T}, \mathcal{D} \rangle$ has already been transformed into the rank equivalent knowledge base $\langle \mathcal{T}^*, \mathcal{D}^* \rangle$ in rank normal form.

In the next lemma, we observe that the inference relation $\vdash_r$ respects the preferential conclusions of $\langle \mathcal{T}, \mathcal{D} \rangle$ w.r.t. the assertions of the form $\top \precsim C$—another desiderata proven for the propositional case by Lehmann and Magidor (1992).

**Lemma 3** *For every concept $C$, $\langle \mathcal{T}, \mathcal{D} \rangle \models_R \top \precsim C$ if and only if $\langle \mathcal{T}, \mathcal{D} \rangle \vdash_r \top \precsim C$.*

Now we need to prove the main result, that is, that our procedure is sound and complete w.r.t. rational closure: We need to prove that, given a knowledge base $\langle \mathcal{T}, \mathcal{D} \rangle$, $\langle \mathcal{T}, \mathcal{D} \rangle \models_R^\cup C \precsim D$ if and only if $\langle \mathcal{T}, \mathcal{D} \rangle \vdash_r C \precsim D$.

The main step consists in checking the correspondence between the ranking function presented in Section 5 for the semantic construction of rational closure ($r_\mathcal{K}$) and the ranking function defined in the above procedure ($r$). First of all, we can check it for the infinite rank.

**Proposition 7** *For every $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$ and every concept $C$, $r_\mathcal{K}(C) = r(C)$.*

Now we can state the main theorem.

**Theorem 5** *Given a knowledge base $\langle \mathcal{T}, \mathcal{D} \rangle$, for every pair of concepts $C, D$, $\langle \mathcal{T}, \mathcal{D} \rangle \models_R^\cup C \precsim D$ iff $\langle \mathcal{T}, \mathcal{D} \rangle \vdash_r C \precsim D$.*

That is, Procedure RationalClosure is correct w.r.t. the definition of rational closure.

We conclude this section with an analysis of the computational complexity of $\vdash_r$, for which we need to analyse procedure RationalClosure on Page 22. Observe that procedure RationalClosure performs at most $n + 2$ (classical) subsumption checks, where $n$ is the number of ranks assigned to elements of $\mathcal{D}$. So the number of subsumption checks performed by procedure RationalClosure is $\mathcal{O}(|\mathcal{D}|)$. Furthermore, we need to execute procedure ComputeRanking on page 19 to obtain the knowledge base $\langle \mathcal{T}^*, \mathcal{D}^* \rangle$ and the sequence $\mathcal{E}_0, \ldots, \mathcal{E}_n$, which are needed as input to procedure RationalClosure. First note that procedure Exceptional on page 18, with $\mathcal{E}$ as input, performs at most $|\mathcal{E}|$ subsumption checks. From this, and an analysis of procedure ComputeRanking, it follows that the number of subsumption checks performed by procedure ComputeRanking is $\mathcal{O}(|\mathcal{D}|^3)$. Since we know that subsumption checking w.r.t. general TBoxes in $\mathcal{ALC}$ is ExpTime-complete (Baader et al., 2007, Chapter 3), we get the following result, which we state without proof.

**Theorem 6** *Computing Rational Closure for $\mathcal{ALC}$ is ExpTime-complete.*

Hence defeasible subsumption checking for general $\mathcal{ALC}$ TBoxes is just as hard as classical subsumption checking.

## 7. Rational Extensions of an ABox

Up until now, we have considered only knowledge bases $\langle \mathcal{T}, \mathcal{D} \rangle$ containing just (classical or defeasible) concept inclusion axioms. In the present section, we consider the extension of the above procedure to knowledge bases containing also an ABox: given information about particular individuals, we want to derive what *presumably* holds about such individuals. Our knowledge base will have a classical ABox, composed of concept and role assertions, but, using the defeasible inclusion axioms in $\mathcal{D}$, we will be able to derive defeasible information about the individuals: we shall indicate with the expression '$a : \tilde{C}$' the conclusion that the individual $a$ *presumably* falls under the concept $C$. A first attempt for this kind of procedure is presented by Casini and Straccia (2010), and a similar version of the following procedure, specified for a entailment relation different from rational closure, appears in another paper (Casini & Straccia, 2014), but they both lack a semantic characterization and the properties of the inference relation are not fully investigated.

We will work with a knowledge base $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$, and, since the procedure for the ABox is built on top of the procedure for the DTBox, from now on we assume that we have already applied Procedure ComputeRanking to the knowledge base $\langle \mathcal{T}, \mathcal{D} \rangle$, that is, we assume that $\langle \mathcal{T}, \mathcal{D} \rangle$ is in rank normal form (hence $\mathcal{D}_\infty$ is empty), and the set $\mathcal{D}$ has already been partitioned into $\mathcal{D}_0, \ldots, \mathcal{D}_n$, for some finite $n$. The idea of the following procedure is to consider each individual named in the ABox to be as typical as possible, that is, to associate with it all the possible defeasible information that is consistent with the rest of the knowledge base. In order to apply the defeasible information locally to each individual, we encode such information using the materializations of the inclusion axioms, i.e., the sets $\overline{\mathcal{D}_i}$ and $\overline{\mathcal{E}_i}$.

Hence, given $\mathcal{D} = \bigcup\{\mathcal{D}_0, \ldots, \mathcal{D}_n\}$, we end up with the sequence of *default concepts* $\Delta = \langle \bigsqcap \overline{\mathcal{E}_0}, \ldots, \bigsqcap \overline{\mathcal{E}_n} \rangle$, as we did in Section 6 at the Step 2 of the Rational Closure procedure. We want to associate with each individual $a \in \mathcal{K}_\mathcal{I}$ (with $\mathcal{K}_\mathcal{I}$ being the set of the individuals named in the ABox) the strongest formula $\bigsqcap \overline{\mathcal{E}_i}$ that is consistent with the knowledge base (remember that $\models \bigsqcap \overline{\mathcal{E}_i} \sqsubseteq \bigsqcap \overline{\mathcal{E}_{i+1}}$ for every $i$, $1 \leq i \leq n$). In such a way we define a new knowledge base $\widetilde{\mathcal{K}} = \langle \mathcal{A}_\mathcal{D}, \mathcal{T} \rangle$, that we call a *rational ABox extension* of the knowledge base $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$.

**Definition 15** [Rational ABox extension] Let $\mathcal{K} = \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$ be a Knowledge Base such that $\langle \mathcal{T}, \mathcal{D} \rangle$ has already been modified by Procedure ComputeRanking in such a way that $\langle \mathcal{T}, \mathcal{D} \rangle$ is in rank normal form and the partition $\{\mathcal{D}_1, \ldots, \mathcal{D}_n\}$ is defined. For $i \leq n$, let $\mathcal{E}_i = \bigcup_{i \leq j \leq n} \mathcal{D}_j$ be as described in Section 6.

A knowledge base $\langle \mathcal{A}_\mathcal{D}, \mathcal{T}, \mathcal{D} \rangle$ is a *rational extension* of $\mathcal{K} = \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$ iff

- $\langle \mathcal{A}_\mathcal{D}, \mathcal{T}, \mathcal{D} \rangle$ is classically consistent and $\mathcal{A} \subseteq \mathcal{A}_\mathcal{D}$.
- For any $a \in \mathcal{K}_\mathcal{I}$, $a : C \in \mathcal{A}_\mathcal{D} \setminus \mathcal{A}$ iff $C = \bigsqcap \overline{\mathcal{E}_i}$ for some $i \leq n$ and for every $\bigsqcap \overline{\mathcal{E}_h}$, $h < i$,

$$\langle \mathcal{T}, \mathcal{A}_\mathcal{D} \cup \{a : \bigsqcap \overline{\mathcal{E}_i}\} \rangle \models \top \sqsubseteq \bot$$

The above definition identifies the extensions of the original ABox $\mathcal{A}$ that associates, with every individual, the defeasible information that is consistent with the rest of the knowledge base. Using such an approach dealing with the individuals, we remain consistent with the idea behind rational closure: the default information still respects the exceptionality

**Procedure** RationalExtension($\mathcal{K}$)

**Input**: KB $\mathcal{K} = \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$, the sequence $\langle \mathcal{E}_0, \ldots, \mathcal{E}_n \rangle$, a sequence $s = \langle a_1, \ldots, a_m \rangle$ of the individuals in $\mathcal{K}_{\mathcal{I}}$

**Output**: KB $\langle \mathcal{A}^s_{\mathcal{D}}, \mathcal{T} \rangle$

$j := 1$;<br>
$\mathcal{A}_{\mathcal{D}} := A$;<br>
**repeat**<br>
    $i := 0$;<br>
    **while** $\langle \mathcal{A}_{\mathcal{D}} \cup \{a_j : \bigsqcap \overline{\mathcal{E}_i}\}, \mathcal{T} \rangle \models \top \sqsubseteq \bot$ *and* $i \leq n$ **do**<br>
        $i := i + 1$;<br>
    **if** $i \leq n$ **then**<br>
        $\mathcal{A}_{\mathcal{D}} := \mathcal{A}_{\mathcal{D}} \cup \{a_j : \bigsqcap \overline{\mathcal{E}_i}\}$;<br>
    $j := j + 1$;<br>
**until** $j := m + 1$;<br>
**return** $\langle \mathcal{A}^s_{\mathcal{D}}, \mathcal{T} \rangle$;

ranking, and we consider each individual to be as typical as possible, preserving the general consistency. Also the semantic characterization that is presented here will confirm that the notion of rational ABox extension is consistent with the basic idea of rational closure, that is, 'pushing' the individuals as low as possible in the ranked model. Still, the main problem is that, since the individuals can be related to each other through roles, the possibility of associating a default concept to an individual is often influenced by the default information associated to other individuals, as shown in the following example.

**Example 3** Consider $\mathcal{K} = \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$, with $\mathcal{A} = \{(a, b) : \mathsf{r}\}$, $\mathcal{T} = \emptyset$ and $\mathcal{D} = \mathcal{D}_0 = \{\top \sqsubsim \mathsf{A} \sqcap \forall \mathsf{r}.\neg \mathsf{A}\}$ (hence we have $\Delta = \langle \bigsqcap \overline{\mathcal{E}_0} \rangle = \langle \mathsf{A} \sqcap \forall \mathsf{r}.\neg \mathsf{A} \rangle$). If we associate $\bigsqcap \overline{\mathcal{E}_0}$ to $a$, we obtain $b : \neg \mathsf{A}$ and we cannot associate $\bigsqcap \overline{\mathcal{E}_0}$ to $b$; on the other hand, if we apply $\bigsqcap \overline{\mathcal{E}_0}$ to $b$, we derive $b : \mathsf{A}$ and we are no longer able to associate $\bigsqcap \overline{\mathcal{E}_0}$ with $a$. Hence, we have *two* possible rational extensions of $\mathcal{K}$.

This implies that, given a knowledge base $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$, even if the rational closure of $\langle \mathcal{T}, \mathcal{D} \rangle$ is always unique, there is the possibility that we have more than one rational ABox extension.

Once we have defined the sequence of default concepts $\Delta$ from $\mathcal{D}$, a simple algorithm to obtain all the possible extensions of a knowledge base $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$, with $\mathcal{K}_{\mathcal{I}}$ the set of the individuals named in $\mathcal{A}$, is shown in Procedure RationalExtension. In the procedure we give as input a linear order $s$ on the individuals in $\mathcal{K}_{\mathcal{I}}$. As seen in Example 3, by associating defeasible information with an individual (see Definition 15), we potentially constrain the information we can associate with another individual. Hence, the order $s$ corresponds in the procedure to a priority in maximising typicality: the lower an individual occurs in $s$, the less constrained we are in assigning to it a minimal rank.

Let $\mathfrak{S}$ be the set containing all such linear orders of the individuals in $\mathcal{K}_{\mathcal{I}}$.

Hence, the procedure returns a knowledge base $\langle \mathcal{A}^s_{\mathcal{D}}, \mathcal{T} \rangle$ for each $s \in \mathfrak{S}$. Now, the following can be proven.

**Proposition 8** *Let $\mathcal{K} = \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$ be a knowledge base, and $\mathcal{K}_{\mathcal{I}}$ the set of the individuals named in $\mathcal{A}$.*
*Given $\mathcal{K}$ and a linear order $s$ of the elements of $\mathcal{K}_{\mathcal{I}}$, Procedure RationalExtension determines a rational ABox extension of $\mathcal{K}$. Contrariwise, every rational ABox extension of $\mathcal{K}$ corresponds to the knowledge base generated by some linear order of the individuals in $\mathcal{K}_{\mathcal{I}}$.*

**Definition 16** [Inference relation $\vdash_r^s$] Given a knowledge base $\mathcal{K} = \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$ and a linear order $s$ on the individuals in $\mathcal{K}_{\mathcal{I}}$, we say that $a : \tilde{C}$ is a *defeasible consequence* of $\mathcal{K} = \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$ w.r.t. the order $s$, noted $\mathcal{K} \vdash_r^s a : \tilde{C}$, if $\langle \mathcal{A}_{\mathcal{D}}^s, \mathcal{T} \rangle \models a : C$, where $\langle \mathcal{A}_{\mathcal{D}}^s, \mathcal{T} \rangle$ is the rational extension generated from $\mathcal{K}$ using the order $s$.

The interesting point of such a entailment relation is that it still satisfies the properties of a *rational* consequence relation, considering their intuitive translation for ABox reasoning.

**Proposition 9** *Given $\mathcal{K}$ and a linear order $s$ of the individuals in $\mathcal{K}$, the inference relation $\vdash_r^s$ satisfies the following properties:*

| | | |
|---|---|---|
| $(REF_{DL})$ | $\langle \mathcal{A}, \mathcal{T}, \Delta \rangle \vdash_r^s a : \tilde{C}$ for every $a : C \in \mathcal{A}$ | *Reflexivity* |
| $(LLE_{DL})$ | $\dfrac{\langle \mathcal{A} \cup \{b : D\}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C} \qquad \models D \equiv E}{\langle \mathcal{A} \cup \{b : E\}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C}}$ | *Left Logical Equivalence* |
| $(RW_{DL})$ | $\dfrac{\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C} \qquad \models C \sqsubseteq D}{\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{D}}$ | *Right Weakening* |
| $(CT_{DL})$ | $\dfrac{\langle \mathcal{A} \cup \{b : D\}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C} \qquad \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s b : \tilde{D}}{\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C}}$ | *Cumulative Transitivity (Cut)* |
| $(CM_{DL})$ | $\dfrac{\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C} \qquad \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s b : \tilde{D}}{\langle \mathcal{A} \cup \{b : D\}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C}}$ | *Cautious Monotonicity* |
| $(OR_{DL})$ | $\dfrac{\langle \mathcal{A} \cup \{b : D\}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C} \qquad \langle \mathcal{A} \cup \{b : E\}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C}}{\langle \mathcal{A} \cup \{b : D \sqcup E\}, \mathcal{T}, \Delta \rangle \vdash_r^s a : \tilde{C}}$ | *Left Disjunction* |
| $(RM_{DL})$ | $\dfrac{\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C} \qquad \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \not\vdash_r^s b : \tilde{\neg D}}{\langle \mathcal{A} \cup \{b : D\}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C}}$ | *Rational Monotonicity* |

**Example 4** We define a DL-variant of the penguin example. Let $\mathcal{K} = \{\mathcal{A}, \mathcal{T}, \mathcal{D}\}$ be a knowledge base where $\mathcal{A} = \{a : \mathsf{P}, b : \mathsf{B}, (a, c) : \mathsf{Hunt}, (b, c) : \mathsf{Hunt}\}$, $\mathcal{T} = \{\mathsf{P} \sqsubseteq \mathsf{B}, \mathsf{I} \sqsubseteq \neg\mathsf{Fi}\}$, $\mathcal{D} = \{\mathsf{B} \sqsubsim \mathsf{F}, \mathsf{P} \sqsubsim \neg\mathsf{F}, \mathsf{B} \sqsubsim \forall\mathsf{Hunt.I}, \mathsf{P} \sqsubsim \forall\mathsf{Hunt.Fi}\}$, where you can read $\mathsf{B}$ as *Bird*, $\mathsf{P}$ as *Penguin*, $\mathsf{F}$ as *Flying*, $\mathsf{I}$ as *Insect*, $\mathsf{Fi}$ as *Fish*, and $\mathsf{Hunt}$ as *hunts*. From $\mathcal{D}$ we obtain the default concepts $\bigsqcap \overline{\mathcal{E}_0} = (\neg\mathsf{B} \sqcup \mathsf{F}) \sqcap (\neg\mathsf{B} \sqcup \forall\mathsf{Hunt.I}) \sqcap (\neg\mathsf{P} \sqcup \neg\mathsf{F}) \sqcap (\neg\mathsf{P} \sqcup \forall\mathsf{Hunt.Fi})$ and $\bigsqcap \overline{\mathcal{E}_1} = (\neg\mathsf{P} \sqcup \neg\mathsf{F}) \sqcap (\neg\mathsf{P} \sqcup \forall\mathsf{Hunt.Fi})$.

Applying our procedure we can identify two possible rational ABox extensions of $\mathcal{K}$: one in which we associate the default concepts first to $a$ and then to $b$, and the second one in which we consider $b$ before $a$. In the former case, we associate with $a$ the default $\bigsqcap \overline{\mathcal{E}_1}$, and we derive that $a$ is a typical penguin that hunts fishes (hence we can conclude $c : Fi$) and does not fly, while, having concluded that $c$ is a fish, we cannot associate anymore $\bigsqcap \overline{\mathcal{E}_0}$ to $b$, and we have to treat $b$ as an atypical bird, and we are not able to associate with $c$ the typical properties of birds, i.e., that it flies and hunts insects. On the other hand, if we consider $b$ before $a$, we associate $\bigsqcap \overline{\mathcal{E}_0}$ to $b$, hence considering $b$ a typical bird that flies and hunts insects, but, being $c$ an insect, we cannot associate with it the concept $\bigsqcap \overline{\mathcal{E}_1}$, and we have to consider $a$ an atypical penguin.

From the point of view of the computational complexity, the decision problem w.r.t. $\vdash^s_r$ has the same complexity as the classical ABox consistency decision problem in $\mathcal{ALC}$ (Donini & Massacci, 2000).

**Proposition 10** *Deciding $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \vdash^s_r a : \tilde{C}$ in $\mathcal{ALC}$ is an ExpTime-complete problem.*

In the presence of multiple rational ABox extensions, we can also define the inference relation $\vdash_r$, a more conservative inference relation independent from any order on the individuals; similar to other *cautious* entailment relations, it is defined via the intersection of all the inference relations $\vdash^s_r$ modeling a rational extension.

**Definition 17** $\vdash_r := \bigcap \{\vdash^s_r \mid s$ is a linear order on the elements of $\mathcal{K}_{\mathcal{I}}\}$.

However, in the presence of multiple ABox extensions there is the possibility that we lose the property of rational monotonicity.

**Proposition 11** *The inference relation $\vdash_r$ does not always satisfy ($RM_{DL}$).*

This is shown by the following example.

**Example 5** Consider the knowledge base $\langle \mathcal{A}, \mathcal{D} \rangle$ s.t. $\mathcal{A} = \{(a, b) : \mathsf{r}\}$ and $\mathcal{D} = \mathcal{D}_0 \cup \mathcal{D}_1$, with $\mathcal{D}_0 = \{\top \mathrel{\sqsubset\!\!\!\sim} \mathsf{A} \sqcap \forall \mathsf{r}.\neg\mathsf{A}, \top \mathrel{\sqsubset\!\!\!\sim} \mathsf{B}\}$ and $\mathcal{D}_1 = \{\neg\mathsf{A} \mathrel{\sqsubset\!\!\!\sim} \neg\mathsf{B}, \neg\forall \mathsf{r}.\neg\mathsf{A} \mathrel{\sqsubset\!\!\!\sim} \mathsf{B}\}$. We can define two sequences on the individuals, $s = \langle a, b \rangle$ and $s' = \langle b, a \rangle$, each of them defining a different rational extension, with $\vdash_r = \vdash^s_r \cap \vdash^{s'}_r$. We have that $\langle \mathcal{A}, \mathcal{D} \rangle \vdash_r a : \mathsf{B}$, since in both the extensions $a : \mathsf{B}$ holds (in $\vdash^s_r$ because of the axiom $\top \mathrel{\sqsubset\!\!\!\sim} \mathsf{B}$ and in $\vdash^{s'}_r$ due to the axiom $\neg\forall \mathsf{r}.\neg\mathsf{A} \mathrel{\sqsubset\!\!\!\sim} \mathsf{B}$) while we have $\langle \mathcal{A}, \mathcal{D} \rangle \not\vdash_r a : \mathsf{A}$, since $\langle \mathcal{A}, \mathcal{D} \rangle \not\vdash^{s'}_r a : \mathsf{A}$. However, $\langle \mathcal{A} \cup \{a : \neg\mathsf{A}\}, \mathcal{D} \rangle \not\vdash_r a : \mathsf{B}$, since $\langle \mathcal{A} \cup \{a : \neg\mathsf{A}\}, \mathcal{D} \rangle \not\vdash^s_r a : \mathsf{B}$.

The computational complexity of $\vdash_r$ is the same as $\vdash^s_r$, i.e., the decision procedure is ExpTime-complete: assuming that the number of individuals named in the ABox is $n$, we have to decide $\vdash^s_r$ for each possible sequences $s$ defined on the $n$ individuals. That is, in the worst case we need to do $n!$ calls of an ExpTime-complete decision procedure, which, again, gives back an ExpTime-complete decision procedure.[6]

In terms of the expected occurrence of multiple rational ABox extensions for real world KBs, there has been initial indication that this would be rare (Casini, Meyer, Moodley, & Varzinczak, 2013b) although it is acknowledged that a more involved investigation is required to be certain of this.

To check whether a knowledge base $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$ has a single rational ABox extension, it is sufficient to associate with each individual in $\mathcal{K}_{\mathcal{I}}$ the strongest $\bigsqcap \overline{\mathcal{E}_i}$ modulo consistency w.r.t $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$, exactly as in the Procedure RationalExtension, but with one difference: we substitute line 5:

6. See e.g., http://lifecs.likai.org/2012/06/better-upper-bound-for-factorial.html.

**while** $\langle \mathcal{A}_\mathcal{D} \cup \{a_j : \bigsqcap \overline{\mathcal{E}_i}\}, \mathcal{T}\rangle \models \top \sqsubseteq \bot$ **do**

with the new line

**while** $\langle \mathcal{A} \cup \{a_j : \bigsqcap \overline{\mathcal{E}_i}\}, \mathcal{T}\rangle \models \top \sqsubseteq \bot$ **do**

That is, we check consistency for each individual without considering the defeasible information we have associated to the others. If the final knowledge base $\langle \mathcal{A}_\mathcal{D}, \mathcal{T}\rangle$ is consistent, it is the only rational ABox extension of $\langle \mathcal{A}, \mathcal{T}, \mathcal{D}\rangle$.

**Proposition 12** *In the presence of a knowledge base $\langle \mathcal{A}, \mathcal{T}, \mathcal{D}\rangle$ that has a single rational ABox extension, checking the uniqueness of the rational ABox extension and, in case, whether $\langle \mathcal{A}, \mathcal{T}, \mathcal{D}\rangle \vdash_r a : \tilde{C}$ is an ExpTime-complete problem in $\mathcal{ALC}$.*

**Example 6** Consider the knowledge base in Example 4, where in $\mathcal{A}$ $(b, c) : \mathsf{Hunt}$ is replaced with $(b, d) : \mathsf{Hunt}$. Then, whatever is the order on the individuals, we obtain the following association between the defeasible information and the individuals: $a : \bigsqcap \overline{\mathcal{E}_1}$, $b : \bigsqcap \overline{\mathcal{E}_0}$, $c : \bigsqcap \overline{\mathcal{E}_0}$, and $d : \bigsqcap \overline{\mathcal{E}_0}$. Using the information in these defaults, we obtain a unique rational ABox extension.

We can also extend the semantic characterisation we have given of rational closure to the ABoxes, so that we can specify a procedure for computing the rational ABox extension we have just defined.

Consider a knowledge base $\langle \mathcal{A}, \mathcal{T}, \mathcal{D}\rangle$. Again, we assume that the tuple $\langle \mathcal{T}, \mathcal{D}\rangle$ is in rank normal form, and $\mathcal{D}$ is partitioned into $\mathcal{D}_1, \ldots, \mathcal{D}_n$. First of all, we can check if it is consistent by using classical reasoning: a knowledge base $\langle \mathcal{A}, \mathcal{T}, \mathcal{D}\rangle$ is consistent if there is a ranked interpretation that satisfies all the assertions in $\mathcal{A}$, all the classical subsumption axioms in $\mathcal{T}$ and all the defeasible subsumption axioms in $\mathcal{D}$. Again, we can reduce this consistency check to classical consistency checks.

**Lemma 4** *A knowledge base $\langle \mathcal{A}, \mathcal{T}, \mathcal{D}\rangle$ in rank normal form is consistent if and only if $\langle \mathcal{A}, \mathcal{T}\rangle \not\models \top \sqsubseteq \bot$.*

Now that we have a method to decide the consistency of a knowledge base $\langle \mathcal{A}, \mathcal{T}, \mathcal{D}\rangle$, we can prove that the consistency of $\langle \mathcal{A}, \mathcal{T}, \mathcal{D}\rangle$ guarantees the existence of an interpretation that extends the rational closure model $R^\mathcal{K}_\cup$ in order to satisfy also the ABox $\mathcal{A}$.

**Lemma 5** *Let $\langle \mathcal{A}, \mathcal{T}, \mathcal{D}\rangle$ be a consistent knowledge base in rank normal form and $R^\mathcal{K}_\cup$ the rational closure model of $\langle \mathcal{T}, \mathcal{D}\rangle$. Then there is an interpretation $R' = \langle \Delta^{R^\mathcal{K}_\cup}, \cdot^{R'}, \prec^{R^\mathcal{K}_\cup}\rangle$ s.t.:*

- *for every atomic concept $A$, $A^{R'} = A^{R^\mathcal{K}_\cup}$;*
- *for every role $r$, $r^{R'} = r^{R^\mathcal{K}_\cup}$;*
- *$R' \Vdash \mathcal{A}$.*

That is, the consistency of a knowledge base $\langle \mathcal{A}, \mathcal{T}, \mathcal{D}\rangle$ implies that there is a characteristic model of the rational closure of $\mathcal{K} = \langle \mathcal{T}, \mathcal{D}\rangle$ (that is, a model that corresponds to $R^\mathcal{K}_\cup$ w.r.t. the interpretation of $\mathcal{K} = \langle \mathcal{T}, \mathcal{D}\rangle$) that also satisfies the ABox $\mathcal{A}$. From this, we can extend the notion of the model of rational closure to a knowledge base including an ABox (again, $\mathcal{K}_{\mathcal{I}}$ is the set of the individuals named in $\mathcal{A}$). In doing this, we define a construction that is based on the semantic approaches presented by Giordano et al. (2013b) and Casini et al. (2013a), in turn extending the minimisation approach by Giordano et al. (2012). We *minimise* the height in the model of the individuals named in the ABox: the lowest the position of the individuals in $\mathcal{K}_{\mathcal{I}}$ inside the model, the more we presume that they behave as typical as possible.

In order to do so, let $\mathfrak{R}_{\mathcal{A}}^{R_{\cup}^{\mathcal{K}}}$ be the set of the models that correspond to $R_{\cup}^{\mathcal{K}}$ w.r.t. the interpretation of $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$ and that also satisfy $\mathcal{A}$. We introduce an order $\leq^{\mathcal{A}}$ on the models in $\mathfrak{R}_{\mathcal{A}}^{R_{\cup}^{\mathcal{K}}}$ such that, given $R, R' \in \mathfrak{R}_{\mathcal{A}}^{R_{\cup}^{\mathcal{K}}}$, $R \leq^{\mathcal{A}} R'$ if and only if $h_R(a^R) \leq h_{R'}(a^{R'})$ for each individual $a$ in $\mathcal{K}_{\mathcal{I}}$. The minimal ABox models of $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$ are the minimal elements of the order $\leq^{\mathcal{A}}$.

**Definition 18** [Minimal ABox model] $R$ is a *minimal ABox model* of a knowledge base $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$ if $R \in \mathfrak{R}_{\mathcal{A}}^{R_{\cup}^{\mathcal{K}}}$ and there is not another model $R' \in \mathfrak{R}_{\mathcal{A}}^{R_{\cup}^{\mathcal{K}}}$ such that $R' \leq^{\mathcal{A}} R$ and $R \not\leq^{\mathcal{A}} R'$.

Let $\mathfrak{M}^{\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle}$ be the set of minimal ABox models of $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$. We indicate with $\mathfrak{M}_h^{\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle}$ the subconcept of $\mathfrak{M}^{\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle}$ composed of the elements of $\mathfrak{M}^{\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle}$ in which each element $a$ of $\mathcal{K}_{\mathcal{I}}$ has a specific height $h(a) = n$.

**Definition 19** We define the entailment relation $\models_h^{\leq}$ as follows: $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \models_h^{\leq} a : C$ iff $\mathscr{M} \Vdash a : C$ for each $\mathscr{M} \in \mathfrak{M}_h^{\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle}$ . We indicate with $\models^{\leq}$ the entailment relation defined by all the minimal ABox models of $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$.

The following proposition describes the correspondence between the inference relations $\vdash_r^s$ and $\vdash_r$ and, respectively, the entailment relations $\models_h^{\leq}$ and $\models^{\leq}$.

**Proposition 13** *Given a knowledge base $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$, each inference relation $\vdash_r^s$ defined by a sequence $s$ on the elements of $\mathcal{K}_{\mathcal{I}}$ corresponds to the entailment relation $\models_h^{\leq}$ for some $h$, and vice versa. The inference relation $\vdash_r$, corresponding to the intersection of all $\vdash_r^s$ generated by $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$, is equivalent to the entailment relation $\models^{\leq}$.*

Next we are going to consider some possible initial optimisations of the ABox queries. Assume we want to know if $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \models^{\leq} a : C$ and we want to draw the safest possible conclusion. In the presence of multiple acceptable extensions, the classical solution is to use a *skeptical* approach, i.e., to use the inference relation $\vdash_r$, corresponding to the intersection of all the inference relations associated to each possible ordering $s$ of the individuals appearing in $\mathcal{A}$.

However, in case of multiple rational extensions, the $\vdash_r$ decision problem becomes more cumbersome from the computational point of view. Observe, though, that the amount of defeasible information associable with an individual $a$ can only be influenced by the individuals related to it by means of a role: it is immediate to see that if there is no role-connection in the ABox between two individuals $a$ and $b$, then the information that is associated with $a$ does not influence the amount of defeasible information that we can associate with $b$.

Hence, we can ease the decisions w.r.t. the ABox introducing the notion of *cluster*, i.e., a set of individuals named in the ABox that are linked by means of a sequence of role connections; see for example the work on islands for classical DLs (Wandelt & Möller, 2012). To do so, given an ABox $\mathcal{A}$, we consider the symmetric and transitive closure of all the role statements in $\mathcal{A}$, and which pairs of individuals named in $\mathcal{A}$ are connected in such a closure.

**Definition 20** [Cluster] Given a knowledge base $\mathcal{K} = \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$, define the relations $\rhd$ and $\approx$ on $\mathcal{K}_{\mathcal{I}}$ as follows: first, $a \rhd b$ if there is an $r$ such that $(a, b) : r \in \mathcal{A}$. Second, let $\approx$ be the reflexive, symmetric, and transitive closure of $\prec$. Then the *cluster* of $a$ (written $[a]_{\approx}$) is the equivalence class of $a$ w.r.t. $\approx$.

In order to check what we can presumably conclude about $a$, it is sufficient to determine $\vdash_r^s$ w.r.t. each sequence $s$ of individuals in $[a]_{\approx}$. Let $\mathcal{A}_a$ be the ABox obtained by restricting $\mathcal{A}$ to the statements containing individuals in $[a]_{\approx}$; the query $a : \tilde{C}$ is decidable using only the knowledge base $\langle \mathcal{A}_a, \mathcal{T}, \mathcal{D} \rangle$.

**Proposition 14** *$\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \vdash_r a : \tilde{C}$ iff $\langle \mathcal{A}_a, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C}$ for every ordering $s$ of the individuals in $\mathcal{A}_a$.*

## 8. Experimental Results

An important question that we ask of Rational Closure, from a practical perspective, is: how much does one pay for the additional expressivity of defeasible subsumption, in terms of practical reasoning performance?

As illustrated in our procedures (Section 6), we have to perform some additional computation over and above the classical decision procedures for DLs. In general, we perform multiple classical entailment checks to perform a single defeasible entailment check.

The question is how much more work are we doing, and we aim to investigate this question in order to provide evidence of the feasibility of adding defeasible features to ontologies.

We present the results of two preliminary evaluations to determine the practical performance of computing Rational Closure for DTBoxes (leaving an investigation of ABoxes for future work). The first uses randomly generated defeasible ontologies of *moderate* size, while the second describes a method for introducing defeasible subsumption into existing real-world ontologies, and uses this data for evaluation.

### 8.1 Purely Synthetic Data

Our first approach is to generate synthetic defeasible ontologies for evaluation. This is agreed to be a sensible preliminary approach to obtain data for evaluation (Bonatti et al., 2011a), since there are no naturally occuring ontologies with explicitly modelled defeasible features.

### 8.2 Data Generation Model

Before we go into detail about the actual method we use to generate artificial ontologies, we discuss how we parametrise our generation methodology. We focus on two broad categories of parameter: *global* and *local* parameters.

Global parameters are those that pertain to the overall metrics of the ontology such as the number of axioms, concepts, roles etc., while local parameters consider the structure of individual axioms (and concept expressions) in the ontology. The latter parameters include factors such as nesting depth of expressions and the length of conjunctions and disjunctions.

We present our selected values for global and local parameters first, and thereafter, discuss the generation of axioms to be used as *entailment queries* in our evaluation.

### Global Parameters

One of the fundamental considerations, which we do not know a priori, is what percentage of the subsumptions in a real world ontology would users make defeasible?

Anecdotal evidence in the literature (Rector, 2004; Hoehndorf, Loebe, Kelso, & Herre, 2007; Schulz, Stenzhorn, Boeker, & Smith, 2009) suggests that the proportion of defeasible vs. strict information in real world ontologies would likely be lower than 25%. However there are various factors which could render this figure unreliable. For example, ontology engineers may have learned over time to avoid representing defeasible information in the TBox part of their ontologies because standardised ontology languages such as the DL-based Web Ontology Language (OWL) (McGuinness, Van Harmelen, et al., 2004) and accompanying editing tools do not support the expression of defeasibility. In fact, the Marine Top Level Ontology (Tzitzikas, Allocca, Bekiari, Marketakis, Fafalios, Doerr, Minadakis, Patkos, & Candela, 2013) is an application ontology that abides by monotonicity as a design practice (Bekiari, Doerr, Allocca, Barde, & Minadakis, 2014, Page 9).

We consider a range of possibilities for the *percentage defeasibility* (ratio of the number of defeasible vs. the number of strict subsumption statements) of an ontology. We consider ten categories for evaluation starting from 10% and increasing in increments of 10 to 100%.

Apart from the proportion of defeasible statements, we conjecture that it seems reasonable to assume that the remaining structure of real world defeasible ontologies might be very similar to that of existing non-defeasible real world ontologies. Therefore, in order to inform the parameterisation of our ontology generation method, it seems prudent to analyse some non-defeasible real world data to gather some metrics to use in our strategy.

We use ontologies from the recently established Manchester OWL Repository (MOWLRep) (Matentzoglu, Tang, Parsia, & Sattler, 2014b) for this purpose. The main motivation behind the establishment of the repository was to address biases in OWL empirical research where experiments are performed on cherry picked data or data lacking sufficient variety. The goal is to provide a platform for sharing high quality data with emphasis on variety for OWL empirical evaluations.

Returning to our current concerns, another parameter to address is the *size* of the ontologies to generate. Again, we can consider a range here. However, the emphasis with this first investigation is to give a very general sense of how defeasible reasoning would perform with ontologies of "reasonable" size, i.e., with "non-toy" examples.

Research in classical DL reasoning optimisation is still grappling with the problem of reasoning with large-scale ontologies (Haarslev & Möller, 2001; Matentzoglu, Tang, Parsia, & Sattler, 2014a), what then to speak of our defeasible reasoning algorithms which have to perform more work over the underlying classical reasoning steps? Therefore, for our purposes of gaining a preliminary insight into the performance of defeasible reasoning, we argue that it is not yet necessary to tackle large-scale ontologies in depth[7].

Despite this, we would like to be somewhat representative of classical real world ontology sizes. Analysing the ontologies in MOWLRep, we found that the median ontology size was around $3,800$ axioms (including non-TBox axioms). When restricting our attention to TBox axioms the median ontology size was $2,200$ axioms. Therefore, we choose to generate ontologies whose maximum sizes are capped at a figure within this range. In our resulting data the maximum ontology sizes were approximately $3,500$ axioms. Such sizes are reasonable and reflect those of numerous real world application ontologies in corpora outside MOWLRep as well (such as the SWEET corpus (Raskin & Pan, 2005)).

In each percentage defeasibility category we would also like to have a minimum size for the generated ontologies. From practical experience working with *application* OWL ontologies (i.e., those not built purely for demonstrational or educational purposes), the minimum sizes we encountered have ranged between $150-250$ axioms. A good example of a small ontology (roughly 150 axioms) that is used for semantic web applications is the Friend of a Friend (FOAF) vocabulary[8]. We therefore choose 150 axioms as our lower bound for ontology size in each percentage defeasibility category.

In summary we generated 35 ontologies in each percentage defeasibility category (a total of 350 ontologies), varying *uniformly* in size between roughly 150 and $3,500$ axioms. We argue that the number 35 is appropriate to give us a good spread of ontology sizes between 150 and $3,500$. The DL $\mathcal{ALC}$ is used to generate each ontology because the theoretical foundation of our algorithms has been explicitly investigated in the context of $\mathcal{ALC}$.

We express the generated ontologies using OWL with OWL/XML syntax so that we can use existing tool support such as parsers, serialisers, reasoners, via the OWL API (Horridge & Bechhofer, 2011).

In order to represent defeasible subsumption in OWL ontologies (it is not a feature of the OWL specification) we "mark" relevant classical subsumption statements in the ontology as defeasible using meta-data constructs in OWL called *OWL annotations*[9]. Such constructs can be associated with specified OWL axioms and, using OWL processing tools such as the Java-based OWL API (Horridge & Bechhofer, 2011), one can programmatically identify the defeasible axioms in an OWL ontology.

Another global parameter considered for ontology generation is *ontology signature size*. Ontology signature refers to the set of concept and role names mentioned explicitly in the ontology. We therefore have to consider the number of concept names and role names to generate per ontology (relative to the number of axioms we wish to generate per ontology). In MOWLRep we found that the number of concept names (respectively role names) per

---

7. It is notable that the size of the ontologies in our dataset cannot be considered large-scale in comparison with some bio-medical ontologies such as those stored in the NCBO BioPortal corpus (Whetzel, Noy, Shah, Alexander, Nyulas, Tudorache, & Musen, 2011). For example, the National Cancer Institute (NCI) thesaurus (Sioutos, Coronado, Haber, Hartel, Shaiu, & Wright, 2007) appears in this corpus and has versions which contain more than $110,000$ axioms. At the same time, the concept hierarchies of most of these large bio-medical ontologies are rather shallow, making them less interesting from the standpoint of reasoning complexity. These ontologies also generally do not make use of all the expressive features available in $\mathcal{ALC}$.
8. `foaf-project.org`
9. `w3.org/TR/owl2-syntax`, Section 10.2

ontology were roughly 40% (respectively 1.5%) of the number of axioms in the ontology. Therefore these values are used in our ontology generation procedure.

Finally, the last global parameter is what we call *DL constructor distribution*. This is basically the proportion of axioms in the ontology which contain a particular DL construct ($\mathcal{ALC}$ construct in our case). That is, for each of the main $\mathcal{ALC}$ concept constructors: negation, disjunction, conjunction, existential and universal role restrictions, we are interested in the percentage of axioms in an ontology that contain this construct (whenever the ontology actually does contain the construct). This is the core global parameter for our generation methodology and, when examining the metrics of MOWLRep ontologies, we found the average values 6.2%, 26.6%, 21.1%, 4.3%, 14% and median values 1.5%, 17.8%, 11.1%, 2.2%, 4.3%, for negation, existential role restrictions, conjunction, disjunction and universal role restrictions, respectively.

### Local Parameters

The structure of individual axioms in our ontologies can be influenced by many parameters. Our task is a little simpler since we only generate subsumption statements in our ontologies. A subsumption statement has a left hand side (LHS) concept expression and a right hand side (RHS) concept expression. The make-up of a subsumption statement is therefore defined by the make-up of its LHS and RHS concept expressions. We focus on two main parameters influencing this make-up: *nesting depth* and *conjunction or disjunction length*.

Nesting depth refers to the number of sub-concept expressions in a given concept expression. For example, the concept name $\mathsf{A}$ has a nesting depth of 1, the expression $\exists \mathsf{R}.\mathsf{A}$ has a nesting depth of 2 (consisting of $\mathsf{A}$ and $\exists \mathsf{R}.\mathsf{A}$) and the expression $\exists \mathsf{R}.(\mathsf{A} \sqcap \mathsf{B})$ has a nesting depth of 4 (consisting of $\mathsf{A}$, $\mathsf{B}$, $(\mathsf{A} \sqcap \mathsf{B})$ and $\exists \mathsf{R}.(\mathsf{A} \sqcap \mathsf{B})$). Syntactic analysis of ontologies in MOWLRep reveals that, on average, the nesting depth of concept expressions in real world ontologies is just 1. That is, the majority of classes in real world ontologies are names.

However, even though the average nesting depth is just 1, we have encountered isolated cases in MOWLRep where this number reaches 188 (and even one ontology where it reaches $1,707$). However, the majority of these larger nesting depths occur in the larger ontologies in MOWLRep (which are much larger than the ontologies in our synthetic dataset), so we opt for a lower maximum nesting depth for our synthetic data.

We omit the strange case of $1,707$ from consideration because it is a single occurrence in the $22,000$ ontologies of MOWLRep. The next highest nesting depth is 188 and the accompanying ontology sizes for such occurrences are in the order of tens of thousands of axioms, whereas we have decided that the ontologies of our dataset should have a maximum of $3,500$ axioms. Therefore, we choose to cap the maximum nesting depth at 19 (one tenth of 188) for our ontologies.

Conjunction and disjunction length refers to the number of conjunct classes or disjunct classes *in a particular level of nesting* for a given concept expression. For example, the *top level* disjunction length of the concept expression $\mathsf{A} \sqcup (\exists \mathsf{R}.\mathsf{B} \sqcup \mathsf{D}) \sqcup \mathsf{C}$ is 3 even though the sub-expression $(\exists \mathsf{R}.\mathsf{B} \sqcup \mathsf{D})$ has a further 2 disjuncts. When examining concept expressions (that actually contain conjunction and disjunction) in the ontologies of MOWLRep, we find that the average conjunction length is around 2 and the average disjunction length is around 2.5. As in the case of nesting depth, the maximum values encountered are much larger.

We encountered a maximum conjunction length of 85 and we therefore choose a maximum of 9 for our data. The maximum disjunction length is 194 and the next highest is 143 but these two cases are the odd ones out in the data (an order of magnitude larger than the remainder of maximum values). We therefore choose the next highest value of 63 (a maximum disjunction length of 6) for our synthetic data.

### Ontology Generation

We feed our selected global and local parameter values into a basic ontology construction procedure. The procedure consists of four main phases (a flow chart of this process is depicted in Figure 3). We give a brief description of each phase here.

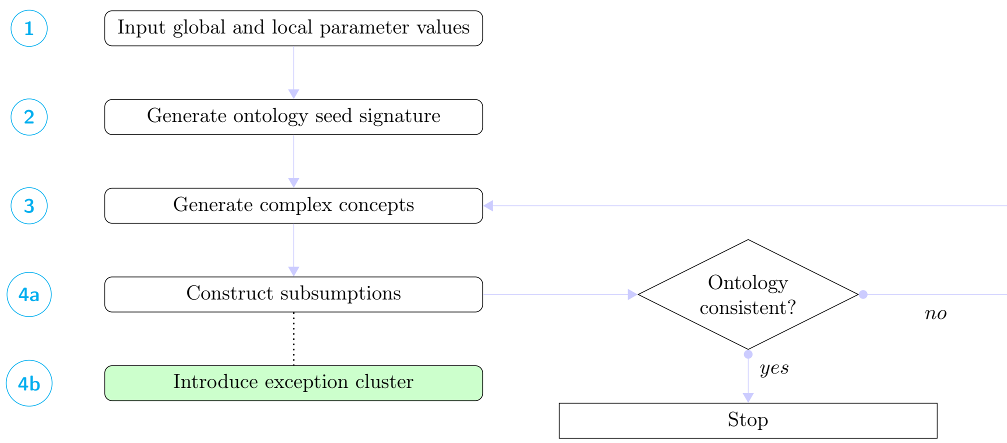


Figure 3: Basic flowchart of artificial ontology generation.

1. *Input global and local parameter values*: the global parameters such as number of axioms to generate, number of concept and role names to use, percentage defeasibility and DL constructor distribution are accepted. Local parameters are also given, namely, the maximum nesting depth and maximum conjunction or disjunction length.

2. *Generate ontology seed signature*: the first main step of the procedure is to generate a set of concept and role names which would be the building blocks for constructing complex concepts and eventually subsumption statements in later steps of the process. If $n$ (respectively $m$) is the number of concept names (respectively role names) to generate, then we generate concept names $\mathsf{A}_1, \ldots, \mathsf{A}_\mathsf{n}$ and role names $\mathsf{R}_1, \ldots, \mathsf{R}_\mathsf{m}$.

We divide the concept names into two equally-sized disjoint sets representing LHS concept names and RHS concept names. This means that the concept names in each respective set are *only* used as either LHS concepts or RHS concepts in subsumptions (and not both). As we shall see in later phases, this is necessary to ensure that we do not have too many syntactic equivalences between concept names in our generated ontology.

3. *Generate complex concepts*: for each concept name in the LHS and RHS sets of Phase 2, we generate three complex concepts containing this concept name. The generated expressions are not divided into LHS and RHS expressions. This results in a total number of concepts that are sufficient to construct all required axioms (recall the requirement to generate 2.5

times more axioms than concept names for the ontology). The provided maximum nesting depth and maximum conjunction and disjunction lengths are used to inform construction.

Since the occurrence of the maximum values for parameters are quite isolated in real world data (MOWLRep), we ensure that this trend is mirrored by our generated data. We use the values obtained for DL constructor distribution to determine the chance of generating a concept expression containing a particular constructor. For example, the chance to generate a concept containing a universal restriction on a role would be around 4%. If we have to generate an existential or universal role restriction then we randomly select from the given set of role names in Phase 2. When we have to generate conjunctions or disjunctions, the conjuncts and disjuncts are randomly selected from the union of the LHS and RHS concept name sets as well as newly introduced complex expressions.

4*a*. *Construct Subsumptions*: analysis of MOWLRep's ontologies reveals that the majority of axioms describe relationships between names (concept names). This is likely because of the use of modern ontology editing tools whose user interfaces are often concept-centric. This leads to specifying axioms predominantly of these forms: $\mathsf{A} \sqsubseteq \mathsf{B}$ or $\mathsf{A} \sqsubseteq \mathsf{C}$ (where $\mathsf{A}$ and $\mathsf{B}$ are concept names and $\mathsf{C}$ is a complex concept). As mentioned earlier we impose a 60% chance of generated axioms of the form $\mathsf{A} \sqsubseteq \mathsf{B}$. A 35% chance is assigned to generate axioms of the form $\mathsf{A} \sqsubseteq \mathsf{C}$. For the remaining 5% of cases we allow the generation of axioms of the form $\mathsf{C} \sqsubseteq \mathsf{D}$ (where both $\mathsf{C}$ and $\mathsf{D}$ are complex expressions). These latter types of axiom (also called general concept axioms) are by far in the minority in real world ontologies. Constructed axioms are added three at a time to the ontology. The reason for this becomes clear in the optional Phase 4*b* which functions in tandem with Phase 4*a*.

4*b*. *Introduce Exception Cluster*: in order to present challenging reasoning cases to our defeasible reasoner, we have to ensure that there are exceptions in our generated ontologies. Our methodology thus far may or may not "organically" introduce exceptions in the ontology, but, to make sure that there are exceptions we assign a small chance to introduce an *exception cluster* into the ontology.

An exception cluster is a set of 3 axioms that represent classic defeasible inheritance example patterns. Consider the student example (Example 2): students generally don't pay taxes ($\mathsf{Student} \precsim \neg(\exists \mathsf{receives}.\mathsf{TaxInvoice})$), but there are specific *types* of students that *do* generally pay taxes, i.e., employed students ($\mathsf{EmployedStudent} \sqsubseteq \mathsf{Student}$, $\mathsf{EmployedStudent} \precsim \exists \mathsf{receives}.\mathsf{TaxInvoice}$).

The general pattern is $C \precsim D$, $E \sqsubseteq C$ and $E \precsim \neg D$. A variant of this pattern is $C \precsim D$, $C \sqcap E \precsim \neg D$. We can also have exceptions-to-exceptions so we can extend both patterns to $C \precsim D$, $E \sqsubseteq C$, $E \precsim \neg D$, $F \sqsubseteq E$, $F \precsim D$ and $C \precsim D$, $C \sqcap E \precsim \neg D$, $C \sqcap E \sqcap F \precsim D$ respectively. The concepts $C$, $D$, $E$ and $F$ are randomly selected from the generated signature and complex expressions.

We impose a 20% chance to introduce an exception cluster each time 3 axioms have been added to the ontology in Phase 4*a*. Even though we only consider up to two levels of exceptions in these patterns, variances in other parameters of ontology generation could increase the overall number of ranks for the resulting defeasible ontologies.

Each time three axioms are generated in Phase 4*a*, we only add them as strict axioms to the ontology. Defeasible axioms are mainly introduced when adding exception clusters (Phase 4*b*). If the desired percentage defeasibility is still not met after the desired ontology size is met, we randomly select the required number of strict axioms from the ontology to be toggled defeasible. Figures 4a and 4b summarise relevant metrics of the generated ontologies.

| Classes | | | Roles | | | Ontology Size | | | Nesting Depth | | Conjuncts | | Disjuncts | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| avg | median | max | avg | median | max | avg | median | max | avg | max | avg | max | avg | max |
| 712 | 713 | 1,376 | 27 | 27 | 52 | 1,799 | 1,800 | 3,503 | 1 | 17 | 2 | 9 | 2 | 4 |

(a) Global and local metrics of generated ontologies.

| Negation (%) | | Conjunction (%) | | Disjunction (%) | | Existential (%) | | Universal (%) | |
|---|---|---|---|---|---|---|---|---|---|
| avg | median | avg | median | avg | median | avg | median | avg | median |
| 1.4 | 1.4 | 17.2 | 17 | 1.9 | 1.9 | 11.6 | 11.6 | 0.13 | 0.1 |

(b) Average percentage of axioms per ontology containing the given $\mathcal{ALC}$ construct.

Figure 4: Relevant metrics and characteristics of the artificial ontologies.

### Entailment Query Generation

In addition to the ontologies we also randomly generated a set of defeasible subsumption statements (entailment queries) for each ontology using terms in their signatures. The number of queries we generated per ontology was 1 percent of the ontology size. In other words, we chose to vary the number of generated queries proportionately according to the size of the ontology.

This was to increase the chances of a generated query set being small enough to be feasible for our experimentation and large/complex enough so that we can gain insights into query performance. All generated queries were stored to file together with their corresponding ontologies.

We argue that the value of 1 percent of ontology size (for the number of queries) is appropriate to give a sufficiently wide range of query times, while still guaranteeing termination of our experiments in reasonable time. For the LHS concept expressions of the entailment queries we randomly selected from the *exceptional* LHS classes of defeasible subsumptions in the ontologies. This is to provide interesting and meaningful queries to our reasoner. If we ask queries with non-exceptional LHSs then defeasible reasoning reduces exactly to classical reasoning (only one classical entailment check is required) and the results would be less interesting for our purposes.

For the RHSs we would like to select expressions that are at least "relevant" to the LHS expression so that it makes sense to actually pose the queries to our reasoner. Therefore, we use the notion of *ontology module* (Grau, Horrocks, Kazakov, & Sattler, 2008) to achieve this. We extract a module (subset) of the ontology that is relevant to the terms in the LHS expression and then collect the terms in this module. The RHS expressions are then randomly generated from these terms.

The generated subsumptions are all defeasible because we are purely interested in the performance of "core" defeasible reasoning. Of course our algorithms themselves do support strict entailment queries but these queries follow from the ontology if and only if they follow via classical entailment from the strict axioms in the ontology. Therefore the performance of such reasoning tasks are more predictable and less interesting to report in our evaluation.

There are, obviously, a variety of ways to generate defeasible entailment queries. However, we conjecture that our strategy represents a sensible first method for an investigation such as ours. Our test data (both ontologies and entailment queries) are publicly available[10].

### 8.3 Experiment Setup

In this section, we give a description of our experimental conditions, the tasks that we execute and the important results we wish to report.

#### Test Setup and Hardware

The first task was to generate the ranking for each of the ontologies in the dataset. We recorded the average time it took to generate a ranking (according to Procedure ComputeRanking in Section 6) for the ontologies of each percentage defeasibility set. The rankings were all stored to file so that they would not have to be recomputed at a later stage.

The second task was to execute Rational Closure (Procedure RationalClosure in Section 6) on the generated set of entailment queries and record the average time to do this per ontology. In terms of optimisations, we use two main techniques for ranking compilation and entailment checking. For ranking compilation the core optimisation is represented by the following relationship between concept unsatisfiability and concept exceptionality:

**Lemma 6** *If a concept $C$ is exceptional w.r.t. a knowledge base $\langle \mathcal{T}, \mathcal{D} \rangle$ then $C$ is unsatisfiable w.r.t. $\mathcal{T} \cup \mathcal{D}'$, where each axiom in $\mathcal{D}'$ is the classical counterpart of a defeasible subsumption in $\mathcal{D}$.*

Lemma 6 states that if a concept is exceptional in a defeasible ontology then it will necessarily be unsatisfiable in the classical translation of the ontology. This result is useful because we can use it to narrow down the search space for identifying exceptional classes in classical ontologies.

Taking the contrapositive of Lemma 6, we obtain the result that if a concept is satisfiable w.r.t. a classical ontology then it is necessarily *not* exceptional w.r.t. any defeasible translation of the ontology. This optimisation can be used to avoid many exceptionality checks while compiling the ranking for an ontology.

Assuming that the ranking is already compiled, we can also optimise the actual procedure for verifying if a given defeasible subsumption is in the rational closure of the KB. This can be done by pruning away axioms from the ranking that are irrelevant to the terms in the query being asked using the notion of module (Grau et al., 2008).

All experiments were executed on an Intel i7 Quad Core machine running Windows 10, with 8GB of memory allocated to the JVM (Java Virtual Machine). Java 1.7 is used with 3GB of memory allocated to the stack for running threads. For loading and analysing the ontologies of our dataset, we use the popular and well-supported Java-based OWL API (Horridge & Bechhofer, 2011).

As we have shown in Section 6, our defeasible reasoning algorithms are built upon classical entailment checks. Thus, we would need to select an existing DL reasoning implementation to perform these classical entailment checks from within our rational closure

10. `github.com/kodymoodley/defeasibleinferenceplatform`

implementation. While running our evaluation with multiple implementations would be interesting for comparison, such an investigation is not necessary to ascertain the price we pay for reasoning with defeasible (in addition to classical) subsumption. We therefore chose to utilise a single DL reasoner for our evaluation. In particular, we would ideally like to use the fastest and most robust implementation.

Consulting the results of the 2014 edition of the OWL Reasoner Evaluation Workshop[11], we identified the top three OWL 2 DL (expressive DL) reasoners for the standard reasoning tasks of: *classification*, *consistency checking* and *satisfiability testing* (in terms of performance and robustness). Robustness was measured as the number of ontologies that were successfully processed in the allotted time. The top reasoners were Konclude[12] (Steigmiller, Liebig, & Glimm, 2014), HermiT[13] (Glimm, Horrocks, Motik, Stoilos, & Wang, 2014), MORe[14] (Romero, Cuenca-Grau, & Horrocks, 2012), Chainsaw[15] (Tsarkov & Palmisano, 2012), FaCT++[16] (Tsarkov & Horrocks, 2006) and TrOWL[17] (Thomas, Pan, & Ren, 2010).

Modern DL reasoners are optimised for classification whereas various other reasoning tasks such as identifying unsatisfiable concept names are usually performed by first classifying the ontology, and then "reading" the relevant information from the results.

Thus, we chose to focus on the reasoners which performed best in OWL 2 DL classification. These were respectively, Konclude, HermiT and MORe. Konclude, unfortunately, does not yet have a direct interface to the OWL API. Therefore, our choice was to select the next best reasoner - HermiT.

### 8.4 Ranking Compilation Results

We view the compilation of the ranking as an "offline" process prior to performing defeasible inference. That is, the ranking should ideally be precompiled and stored to file whenever there is a stable version of the ontology. When reasoning needs to be conducted then the ranking is loaded and entailment queries can be posed (the ranking should not be computed as part of every entailment query).

That being said, the ranking times we obtained for our data seem very reasonable considering that we have implemented only one optimisation for Procedure ComputeRanking. As a point of reference, the average ranking times we observed in our data are comparable to the average times to compute all justifications for an entailment in the BioPortal corpus of ontologies (Horridge, 2011, Figure 6.1, page 99). The percentile plot in Figure 5 gives a summary of the ranking times.

Percentile plots are chosen to represent the data because it gives a better general picture of the performance for the majority of the data, and it also helps to reveal the outlier cases more clearly. For example, if we obtain a value of 5 seconds for the 90th percentile (P90) then it means that 90% of ontologies in the dataset could be ranked in 5 seconds or less.

11. `dl.kr.org/ore2014/results.html`
12. `derivo.de/produkte/konclude.html`
13. `hermit-reasoner.com`
14. `cs.ox.ac.uk/isg/tools/MORe`
15. `chainsaw.sourceforge.net`
16. `owl.man.ac.uk/factplusplus`
17. `trowl.org`

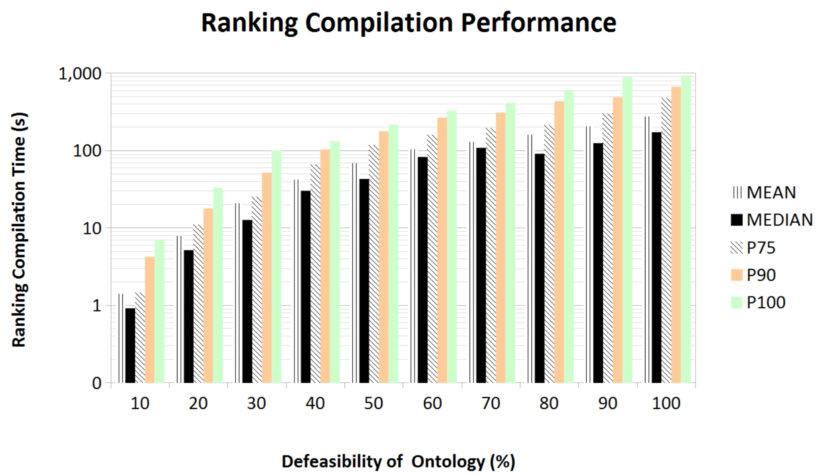


Figure 5: Average time to compute the ranking for the artificial ontologies.

By this definition we note that the 50th percentile is actually the median value for a given dataset and P100 is the maximum value obtained.

Looking at the percentile plot of the ranking times in Figure 5, it seems that ranking compilation gets harder in a fairly uniform manner as the percentage of defeasibility increases. This behaviour is to be expected since as percentage defeasibility increases, the proportion of subsumptions that have unsatisfiable LHS concept expressions potentially also increases. Recall that for the ranking procedure, we have to perform an exceptionality check w.r.t. the ontology for each of these concept expressions.

In addition, as we anticipated, the general trend is that ranking times increase with the number of ranks (also called the *ranking size* or the length of the exception-to-exception chain). Figure 6 illustrates this trend.

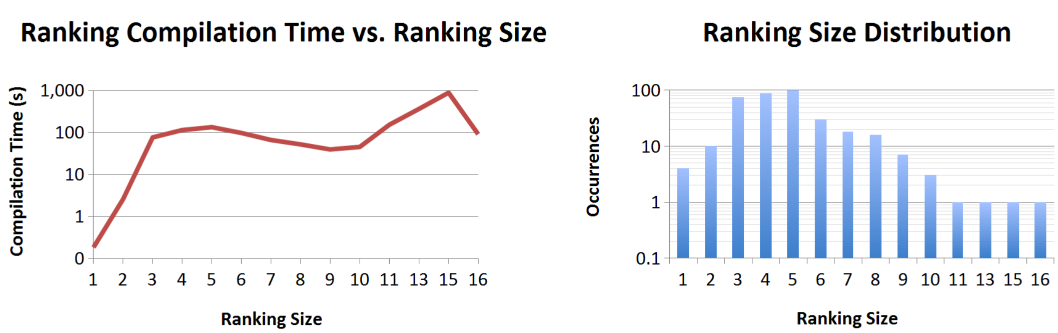


(a) Average ranking time versus the number of ranks in the ranking.

(b) Ranking sizes encountered together with their frequencies.

Figure 6: Ranking size vs. ranking compilation performance. The Y-axis in Figure 6b denotes the number of ontologies in our data that have the indicated ranking size.

We observe that there is a dip in the curve between the ranking sizes of 5 and 10 and also between 15 and 16. The reason for these breaks in the trend is that these portions of the data coincide with brief declines in the percentage defeasibility of ontologies (percentage defeasibility is the other major factor influencing ranking compilation time).

Another important factor is the number of times we have to recurse on the ranking procedure to filter out the hidden strict subsumptions. It is sensible to anticipate that when this recursion factor increases, our ranking times will also increase. This is confirmed in Figure 7.

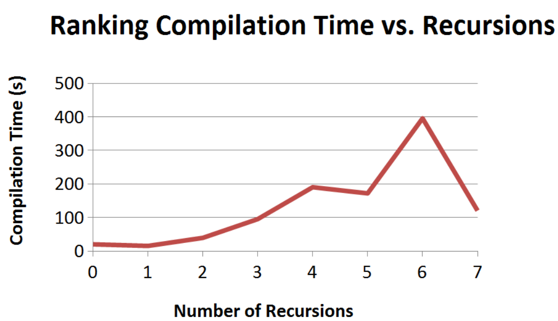


(a) Ranking time vs recursion count.

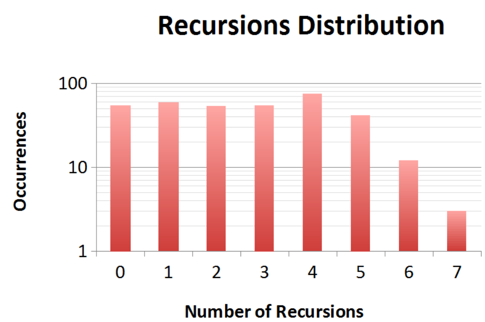


(b) Recursion count frequencies.

Figure 7: Recursion counts vs. ranking compilation performance. The Y-axis in Figure 7b denotes the number of ontologies in our data that have the indicated recursion count.

We have two dips in the curve of Figure 7a. One between recursion counts 4 and 5 and another between 6 and 7. These dips also coincide with declines in percentage defeasibility (from 72% to 64% and 83% to 67% respectively).

It must also be mentioned that the reliability of the curve shape in Figure 7a is greater between recursion counts 0 and 5. These are the most frequent counts found in the data (see Figure 7b) and thus the corresponding average values for the ranking time are more reliable in this range. The same can be said of the ranking size range between 3 and 5 for Figures 6a and 6b. We conclude this section with a summary of average metrics pertaining to the ranking compilation over the entire dataset (see Figure 8).

| Defeasible Axioms | Ranking Size | Hidden Strict. | Rank 0 Size | Rank Size | Exception. Checks | Exceptional LHSs | Unsat. LHSs | Ranking Time (s) |
|---|---|---|---|---|---|---|---|---|
| 991.2 | 4.8 | 500.6 | 312.9 | 98.5 | 2879.2 | 89.4 | 309.8 | 101.4 |

Figure 8: Average metrics pertaining to the ranking compilation per ontology. From left to right: number of defeasible axioms, ranking size, number of hidden strict subsumptions, size of the first rank (containing the non-exceptional defeasible axioms), number of axioms in a general rank, number of exceptionality checks to compute a ranking, number of exceptional LHS concepts of defeasible subsumptions, number of unsatisfiable LHS concepts of defeasible subsumptions and time to compute a ranking.

### 8.5 Entailment Checking Results

Recall that the main goal is to get a general idea of the practical performance of rational closure as well as insights into where the main bottlenecks lie for this procedure.

That said, the results are very encouraging even though the performance degrades significantly (but not drastically) as the percentage defeasibility of the ontologies increase. The overall performance is illustrated in Figure 9.

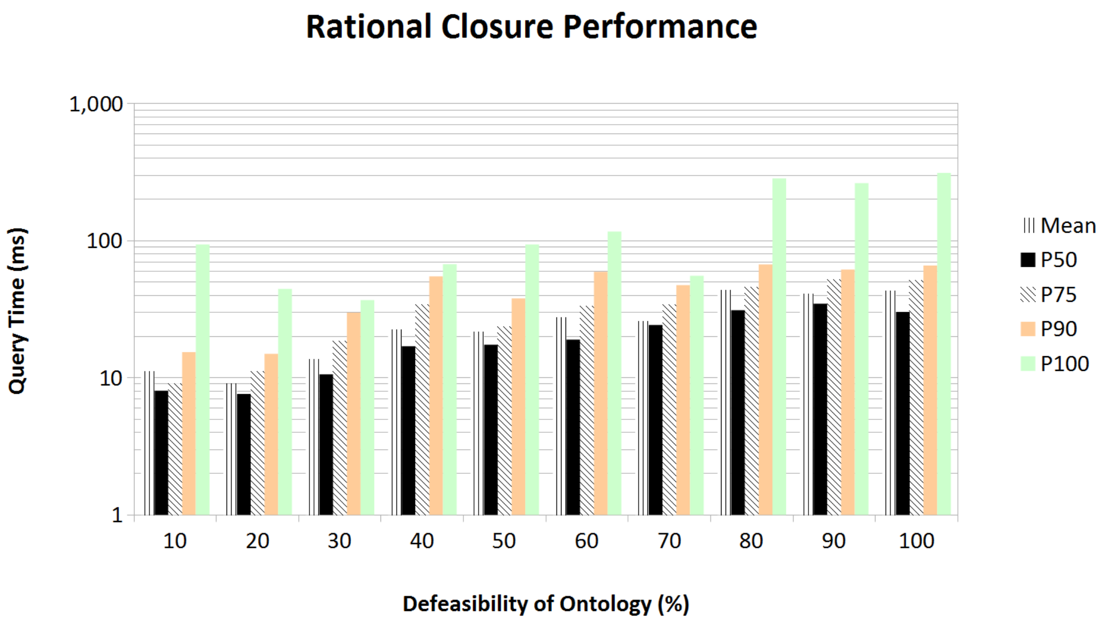


Figure 9: Average entailment checking times for Rational closure.

It is noteworthy that the average defeasible inference times using Rational closure range between just 11ms and 43ms across the dataset (the maximum average time taken to compute an inference for any individual query in the dataset was 313ms).

The data places us in a position to give a preliminary answer to one of our questions at the start of this chapter (i.e., how much more intensive, on average, is defeasible reasoning than classical reasoning?). We plot the average number of classical entailment checks we use per defeasible entailment check for Rational closure in Figure 10.
As we can see the graph depicts how many classical entailment checks it takes (on average) to compute a single defeasible entailment check. It is interesting that this value stays fairly consistent around the 3.5 mark for Rational closure across the different percentage defeasibilities. Since classical entailment checks are, by far, the most computationally intensive components of our procedures for Rational closure, we can make the generalisation that Rational closure would likely take roughly 3.5 times as long as classical inference for $\mathcal{ALC}$ (for real world defeasible ontologies that have sizes represented in our dataset).

In terms of performance bottlenecks, one presumption is that the rank of the antecedent concept of the query being posed would be the major performance indicator for Rational Closure. The graph depicted in Figure 11a does not support this hypothesis.

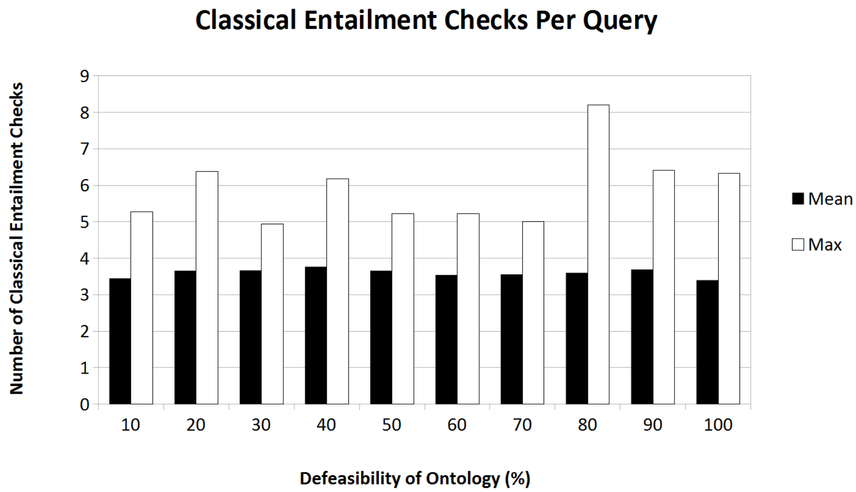


Figure 10: Average and maximum number of classical entailment checks per defeasible entailment check using Rational Closure.

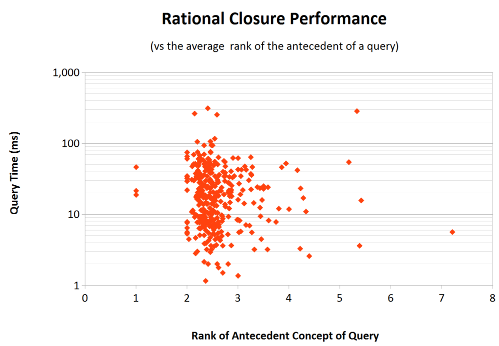


(a) Query time vs. rank of the antecedent.

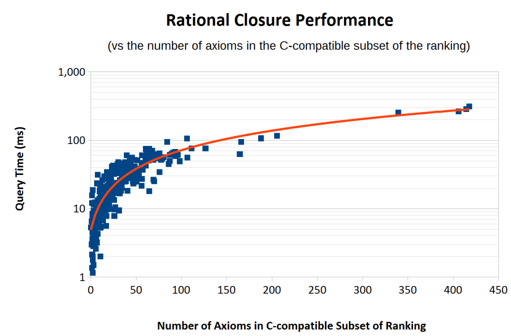


(b) Query time vs. size of $C$-compatible subset.

Figure 11: Potential and actual main bottlenecks for Rational closure.

Judging from the data points in the graph there seems to be no consistent increase in query times as the average rank of the antecedent increases. What we can observe is that the average rank of query antecedents in our dataset lies predominantly between 2 and 3. Nevertheless, there is considerable variance in query time for this range of average antecedent ranks, from roughly 2ms to 150ms (Figure 11a depicts a logarithmic scale). Therefore, there must be another variable which is contributing more to Rational closure query time.

As mentioned earlier, the most computationally intensive components of the Rational closure procedure are its classical entailment checks (each one is an EXPTIME-COMPLETE problem in the worst case for $\mathcal{ALC}$). Figure 10 has illustrated that the number of classical entailment checks does not vary much around 3.5 (not enough to warrant the large variance in query times between the antecedent ranks of 2 and 3).

Therefore, it stands to reason that it is likely that the individual classical entailment checks themselves are taking longer than usual to compute for the hard cases. In other words, we shift to the suspicion that we are inheriting the main bottleneck for classical entailment checking - well known to be ontology size (Sazonau, Sattler, & Brown, 2014).

This suspicion is actually confirmed by the correlation shown in Figure 11b. I.e., in this graph we can see that as the number of axioms in our $C$-compatible subset of the ranking increases, the query times also increase. The increase is quite dramatic until the $C$-compatible subset size is between 50 and 100 (the scale is logarithmic), thereafter the query times actually start to taper but still increase (although the number of occurrences of $C$-compatible subset sizes above this range also decreases dramatically). It is worth mentioning, then, that ontology size (the number of defeasible axioms in the ontology) will always be a significant factor on performance for all our defeasible reasoning algorithms.

### 8.6 Modified Real-world Data

In this section, we take a step further than using purely synthetic ontologies. We describe a principled way of introducing defeasible subsumption into real-world ontologies. We then perform an evaluation of the performance of Rational Closure on the resulting data.

Previously, in terms of data for the evaluation of defeasible reasoning performance, the norm has been the use of automatically generated ontologies with defeasible features. The most notable attempt at a benchmark of synthetic defeasible ontologies is LoDEN[18]. Indeed, we have also used synthetic data in Section 8.1 as a preliminary indicator of performance. Naturally, there are obvious shortcomings with this kind of approach, such as possible biases in the ontology generation methodology.

Here we choose a middle-ground approach, taking advantage of the rich set of (classical) OWL ontologies that we have on the Web in various repositories and corpora. The basic idea of our approach is to modify selected subsumptions in these ontologies to be *defeasible* subsumptions, making these ontologies useful as data to evaluate our defeasible reasoner.

Of course, this has to be done with care to generate cases which are challenging for the reasoner. For example, we need to ensure that there are cases where there are more than one rank in the ranking of the ontology (see Procedure ComputeRanking).

#### Experiment Setup

In this section, we discuss the data curation for our initial set of unmodified ontologies, the methodology for introducing defeasible subsumption into these ontologies, and our experimental conditions.

**Non-defeasible Dataset:** For our initial data, we sample some classical OWL ontologies which we can later pass through our procedure for the introduction of defeasible features. The natural choice is to select the same data that is traditionally used to evaluate the performance of existing *classical* DL reasoners.

However, even in such a setting, there is no precise concensus on what data to use. Therefore, data is generally curated manually by choosing "well-known" ontologies or corpora from which to sample, or arbitrarily selecting from the well-known corpora on the web.

18. `loden.fisica.unina.it`

*Choice of corpora:* While there are existing ontology benchmarks such as LUBM (Guo, Pan, & Heflin, 2005) and its extension UOBM (Ma, Yang, Qiu, Xie, Pan, & Liu, 2006), it was pointed out that there are shortcomings in manual selection of ontologies and ontology corpora for evaluation (Matentzoglu, Bail, & Parsia, 2013). In particular, the main limitation with such selection procedures is that they result in datasets lacking sufficient *variety*.

Thus the results of evaluations can be heavily skewed or biased towards the particular benchmarks being used. As mentioned in Section 8.1, the Manchester OWL Repository (Matentzoglu et al., 2014a) is an effort to address this issue. The current version of the repository contains three core datasets, namely versions of NCBO Bioportal[19] (Whetzel et al., 2011), The Oxford Ontology Library (OOL)[20] and MOWLCorp[21] (Matentzoglu et al., 2013).

While Bioportal and OOL are already established ontology corpora that are actively used in DL reasoner evaluations, MOWLCorp is a recent gathering of ontologies through sophisticated web crawls and filtration techniques.

We obtain a recent snapshot of the Manchester OWL Repository as the base dataset for our evaluation. There are 344, 793 and 20,996 ontologies in the Bioportal, OOL and MOWLCorp corpora respectively.

*Filtration Process and Choice of DL Reasoner:* For loading and analysing the ontologies of our dataset, we use the popular and well-supported Java-based OWL API (Horridge & Bechhofer, 2011). Just as in the case for the artificial data, we use the classical reasoning implementation - HermiT - to perform our classical entailment checks.

Given our choice of tools for manipulating and reasoning with the ontologies in our dataset, we extracted the ontologies that could be loaded and parsed by the OWL API (each within an allotted 40 minutes). The resulting ontologies were then tested to determine if they were classifiable by HermiT within an additional 40 minutes each. Those ontologies which did not pass this test were also removed from the data.

Ontology size has been shown to be the dominant factor for classical reasoning performance (Sazonau et al., 2014), therefore we remove from consideration those ontologies which have less than 150 logical axioms (in keeping with our artificial data which had the same minimum ontology sizes). This is because there is no reason to believe that smaller ontologies than this would be any more interesting or challenging for our reasoner.

Finally, we stripped the ontologies of ABox data because we are only interested in testing reasoning performance for DTBoxes. This leaves us with 252, 440 and 2335 ontologies in Bioportal, OOL and MOWLCorp respectively.

**Defeasible Dataset:** In this section, we describe a systematic technique to introduce defeasible subsumptions into the ontologies of our dataset, thereby making them amenable to defeasible reasoning evaluation.

*Methodology:* Our approach hinges upon Lemma 6 which states that if a concept is exceptional in a defeasible ontology then it will necessarily be unsatisfiable in the classical translation of the ontology. This result is useful because we can use it to *approximate* or identify *possible* exceptional classes in classical ontologies.

19. `bioportal.bioontology.org`
20. `cs.ox.ac.uk/isg/ontologies`
21. `mowlrepo.cs.manchester.ac.uk/datasets/mowlcorp`

Again, the contrapositive of Lemma 6 states that if a concept is satisfiable w.r.t. a classical ontology then it necessarily cannot be exceptional w.r.t. any defeasible translation of the ontology. Therefore, we can eliminate ontologies from our dataset without LHS-classes of subsumptions that are unsatisfiable, because these ontologies could never contain exceptions of the kind we define in this work.

The following definition, which is a generalisation of standard incoherence to axioms with complex left hand side (LHS) concepts, helps us to define these cases:

**Definition 21** A classical TBox $\mathcal{T}$ is *LHS-coherent* if each $C \sqsubseteq D \in \mathcal{T}$ is s.t. $\mathcal{T} \not\models C \sqsubseteq \bot$. $\mathcal{T}$ is *LHS-incoherent* if it is not LHS-coherent.

Eliminating all ontologies from our dataset that are LHS-coherent leaves us with 11, 46 and 77 ontologies in the Bioportal, OOL and MOWLCorp corpora respectively. Figure 13 provides some average properties of the ontologies in our dataset.

| Corpus | Classes | | | Roles | | | TBox size | | | RBox Size | | | Nested Classes | | Conjuncts | | Disjuncts | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | avg | median | max | avg | median | max | avg | median | max | avg | median | max | max | avg | max | avg | max | avg |
| Bioportal | 17992 | 422 | 187515 | 50 | 33 | 152 | 41309 | 754 | 439208 | 30 | 17 | 149 | 28 | 1 | 10 | 2 | 5 | 2.1 |
| OOL | 22203 | 16306 | 89926 | 55 | 44 | 194 | 41389 | 32928 | 142996 | 27 | 9 | 123 | 136 | 1 | 68 | 2 | 16 | 2.5 |
| MOWLCorp | 4621 | 216 | 89926 | 466 | 13.5 | 16586 | 8719 | 691 | 137021 | 214 | 5 | 7749 | 34 | 1 | 17 | 2 | 16 | 2.7 |

Figure 12: Ontology metrics for the LHS-incoherent cases in the dataset.

In total we have 134 ontologies for our performance evaluation. The task is to "relax" some of the subsumptions of our ontologies to be defeasible. The obvious naïve approach to introducing defeasibility would be to convert *all* subsumptions to defeasible ones. This is perhaps not a likely general design methodology of defeasible-ontology engineers in practice.

The other extreme would be to develop an approach to identify the *minimal* (for some defined notion of minimality) amount of defeasibility to introduce into the ontology in order to successfully "cater for all the exceptions". The latter approach would be interesting, and we are currently investigating such an approach; however, we propose that a reasonable approximation of such a procedure yields meaningful data for performance evaluation.

We illustrate the problem by means of an example:

**Example 7** Consider the following TBox $\mathcal{T}$ about different types of mechanics (Casini, Meyer, Moodley, Sattler, & Varzinczak, 2015):

1. $\mathsf{Mechanic} \sqsubseteq \exists\mathsf{hasWorkshop}.\top$,
2. $\mathsf{Mechanic} \sqsubseteq \exists\mathsf{hasSpecialisation}.\top$,
3. $\mathsf{MobileMechanic} \sqcup \mathsf{GeneralMechanic} \sqcup \mathsf{CarMechanic} \sqsubseteq \mathsf{Mechanic}$,
4. $\mathsf{MobileMechanic} \sqsubseteq \neg\exists\mathsf{hasWorkshop}.\top$,
5. $\mathsf{MobileMechanic} \sqcap \neg\exists\mathsf{status}.\mathsf{OnStandBy} \sqsubseteq \exists\mathsf{hasWorkshop}.\top$,
6. $\mathsf{GeneralMechanic} \sqsubseteq \neg\exists\mathsf{hasSpecialisation}.\top$,
7. $\mathsf{CarMechanic} \sqsubseteq \exists\mathsf{hasSpecialisation}.\mathsf{Car}$

The classes $\mathsf{MobileMechanic}$, $\mathsf{GeneralMechanic}$ and the concept expression $\mathsf{MobileMechanic} \sqcap \neg\exists\mathsf{status}.\mathsf{OnStandBy}$ are unsatisfiable w.r.t. $\mathcal{T}$. An intuitive analysis of $\mathcal{T}$ tells us that the ontology engineer may have intended to model that mechanics *usually* have a workshop ($\mathsf{Mechanic} \mathrel{\sqsubseteq\!\!\!\sim} \exists\mathsf{hasWorkshop}.\top$) and *usually* specialise in certain types of equipment that they repair ($\mathsf{Mechanic} \mathrel{\sqsubseteq\!\!\!\sim} \exists\mathsf{hasSpecialisation}.\top$).

This translation of Axioms 1 and 2 in Example 7, is a minimal and intuitive way to introduce defeasibility into $\mathcal{T}$, catering for exceptional types of mechanic - i.e., mobile and general mechanics.

However, we also have an exceptional type of mobile mechanic in $\mathcal{T}$ (an "exception-to-an-exception"). That is, mobile mechanics who are no longer "on standby" or "on call" ($\mathsf{MobileMechanic} \sqcap \neg\exists\mathsf{status}.\mathsf{OnStandBy}$). These mechanics would then be assigned a workshop for their repair tasks.

To cater for such mechanics we would have to relax Axiom 4 as well of Example 7 to express that mobile mechanics *usually* don't have a workshop ($\mathsf{MobileMechanic} \mathrel{\sqsubseteq\!\!\!\sim} \neg\exists\mathsf{hasWorkshop}.\top$).

We now define a general *defeasible translation function* (DTF) for converting classical subsumptions to defeasible subsumptions in classical ontologies.

**Definition 22** (DTF) Let $\mathcal{T}$ be a set of classical subsumptions of the form $C \sqsubseteq D$, then $\mathscr{F} : \mathcal{T} \to \{C \mathrel{\sqsubseteq\!\!\!\sim} D \mid C \sqsubseteq D \in \mathcal{T}\} \cup \mathcal{T}$ is a *DTF* for $\mathcal{T}$.

We also have to formalise what we mean when a particular DTF "caters for all exceptions" in the TBox. We call such a function a *safe* DTF.

**Definition 23** (safe DTF) Let $\mathcal{T}$ be a set of classical subsumptions, let $\mathscr{F}$ be a DTF for $\mathcal{T}$ and let $\mathscr{D}$ be the special DTF that translates *all* subsumptions in $\mathcal{T}$ to defeasible ones. Then, $\mathscr{F}$ is a *safe DTF* for $\mathcal{T}$ if $C$ is totally (resp. normally) exceptional w.r.t. $\mathscr{D}(\mathcal{T})$ if $C$ is totally (resp. normally) exceptional w.r.t. $\mathscr{F}(\mathcal{T})$, for each $C \sqsubseteq D \in \mathcal{T}$.

Intuitively speaking, Definition 23 describes a DTF that preserves the type of the exception exhibited in the *full* defeasible translation of the classical ontology (all subsumptions converted to defeasible). That is, if there is a *normal exception* in the full defeasible translation, then any DTF that translates fewer axioms should preserve this normal exception (It should not change to a *total* exception in the translated ontology.).

We define a safe DTF placing a small upper bound on the subset of axioms to relax using the well-known notion of *justification* (Horridge, 2011). A justification for some entailment $\alpha$ of an ontology is a minimal (w.r.t. set inclusion) subset of the ontology that entails $\alpha$.

If we compute the justifications for $\mathcal{T} \models \mathsf{MobileMechanic} \sqsubseteq \bot$ (the concise reasons for $\mathsf{MobileMechanic}$ being unsatisfiable and possibly exceptional) we obtain a single justification $\{1, 3, 4\}$. Relaxing these axioms is sufficient for catering for mobile mechanics (in fact, it is only necessary to relax Axiom 1 as mentioned earlier). Simlarly, we arrive at $\{2, 3, 6\}$ to cater for general mechanics and $\{4, 5\}$ for mobile mechanics no longer on call.

The basic idea is thus to take the union of the justifications for the unsatisfiable LHS-classes and relax these axioms to defeasible ones. We obtain that $\{1, 2, 3, 4, 5, 6\}$ should be relaxed in Example 7, which is admittedly a large proportion of our TBox. However, we have discovered that this proportion is much smaller in practice (Casini et al., 2015).

While computing all justifications has been shown to be feasible in general on real-world ontologies, *black-box* (reasoner-independent) procedures are known to be exponential in the worst case (Horridge, 2011). To reduce the search space for the justifications, we extract the *star locality based module* (Sattler, Schneider, & Zakharyaschev, 2009) for the ontology in question w.r.t. the set of unsatisfiable LHS-classes. A *module* of an ontology w.r.t. a signature (set of terms from the ontology) is a (in this case ideally small) subset of the ontology that preserves the meaning of the terms in the signature (Grau et al., 2008). We specifically choose star locality based modules because of two key properties: (i) they preserve all justifications in the ontology for all entailments (or axioms) that can be constructed with the given signature (*depleting* property (Sattler et al., 2009, Section 3)), and (ii) they are smaller in size relative to other modules which have the depleting property. The pseudocode of our method is given in Procedure relaxSubsumption.

**Procedure** relaxSubsumption($\mathcal{O}$, $\mathcal{C}$)

**Input**: LHS-incoherent TBox $\mathcal{O}$, $\mathcal{C} = \{C \mid (C \sqsubseteq D \in \mathcal{O} \text{ for some } D) \wedge (\mathcal{O} \models C \sqsubseteq \bot)\}$
**Output**: Defeasible ontology $\langle \mathcal{T}, \mathcal{D} \rangle$

$\mathcal{T} := \emptyset$; $\mathcal{D} := \emptyset$;
$\mathcal{M} := \mathit{extractStarModule}(\mathcal{O}, \mathit{sig}(\mathcal{C}))$; $\mathcal{T} := \mathcal{O} \backslash \mathcal{M}$;
**foreach** $X \sqsubseteq Y \in \mathcal{M}$ **do**
  $\mathcal{D} := \mathcal{D} \cup \{X \mathrel{\underset{\sim}{\sqsubset}} Y\}$;
**return** $\langle \mathcal{T}, \mathcal{D} \rangle$;

Finally, it can be shown that our procedure defines a safe DTF for knowledge bases.

**Lemma 7** *(safety of our DTF) Let $\mathscr{F}$ be the DTF defined by Procedure relaxSubsumption, and $\mathcal{O}$ a classical TBox, then $\mathscr{F}$ is a safe DTF for $\mathcal{O}$.*

*Discussion:* There are two main issues with the procedure we have presented for introducing defeasibility into OWL ontologies: (i) minimality of modification to the original ontology and (ii) the representative quality of the resulting defeasible ontology as something that might be built by a ontology engineer with access to defeasible features.

While (i) and (ii) would be useful goals for a methodology automating the introduction of defeasible features into OWL ontologies, our approach does not yet meet such desiderata. It is clear that the minimal axioms to relax in Example 7 would be $\{1, 2, 4\}$, yet we relax $\{3, 5, 6\}$ as well.

The resulting ontology should also ideally resemble a naturally occurring ontology with defeasible features introduced where explicitly needed by the ontology engineer. For instance, in Example 7, it is not intuitive to relax $\mathsf{MobileMechanic} \sqsubseteq \mathsf{Mechanic}$ (all mobile mechanics are mechanics) to $\mathsf{MobileMechanic} \mathrel{\underset{\sim}{\sqsubset}} \mathsf{Mechanic}$ (*typical* mobile mechanics are mechanics). Such constraints should ideally remain strict.

Furthermore, a critical observation is that logical incoherence in classical ontologies may be caused by erroneous modelling. In ontology development tools, large emphasis has been placed on debugging incoherence by making modifications to the ontology to remove the "unwanted" entailments such as $C \sqsubseteq \bot$. This is likely to have prevented many developers publishing incoherent ontologies.

Given the above main shortcomings of our approach, we do not argue that ours is the ideal methodology. Rather, we hope that it serves as a stepping stone from purely synthetic approaches to investigate and develop more suitable methodologies.

**Test Setup and Hardware:** Our setup, methodologies and design choices for the experimental evaluation can be summarised as follows:
*Data summary:* The input data for our experiments are 134 LHS-incoherent ontologies (curated as described in Section 8.6) from the Manchester OWL Repository. The ontologies are divided across three corpora: 11 , 46 and 77 in Bioportal, OOL and MOWLCorp respectively. Figure 13 provides some average properties of the ontologies in our dataset.

| Corpus | Classes | | | Roles | | | TBox size | | | RBox Size | | | Nested Classes | | Conjuncts | | Disjuncts | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | avg | median | max | avg | median | max | avg | median | max | avg | median | max | max | avg | max | avg | max | avg |
| **Bioportal** | 17992 | 422 | 187515 | 50 | 33 | 152 | 41309 | 754 | 439208 | 30 | 17 | 149 | 28 | 1 | 10 | 2 | 5 | 2.1 |
| **OOL** | 22203 | 16306 | 89926 | 55 | 44 | 194 | 41389 | 32928 | 142996 | 27 | 9 | 123 | 136 | 1 | 68 | 2 | 16 | 2.5 |
| **MOWLCorp** | 4621 | 216 | 89926 | 466 | 13.5 | 16586 | 8719 | 691 | 137021 | 214 | 5 | 7749 | 34 | 1 | 17 | 2 | 16 | 2.7 |

Figure 13: Ontology metrics for the LHS-incoherent cases in the dataset.

We also give an illustration in Figure 14 of how much defeasibility our methodology has introduced in to the curated ontologies. The average ratio of defeasible to strict axioms in each ontology is 6%, the median being 1%, the minimum ratio being 0.01% and the maximum being 98%.

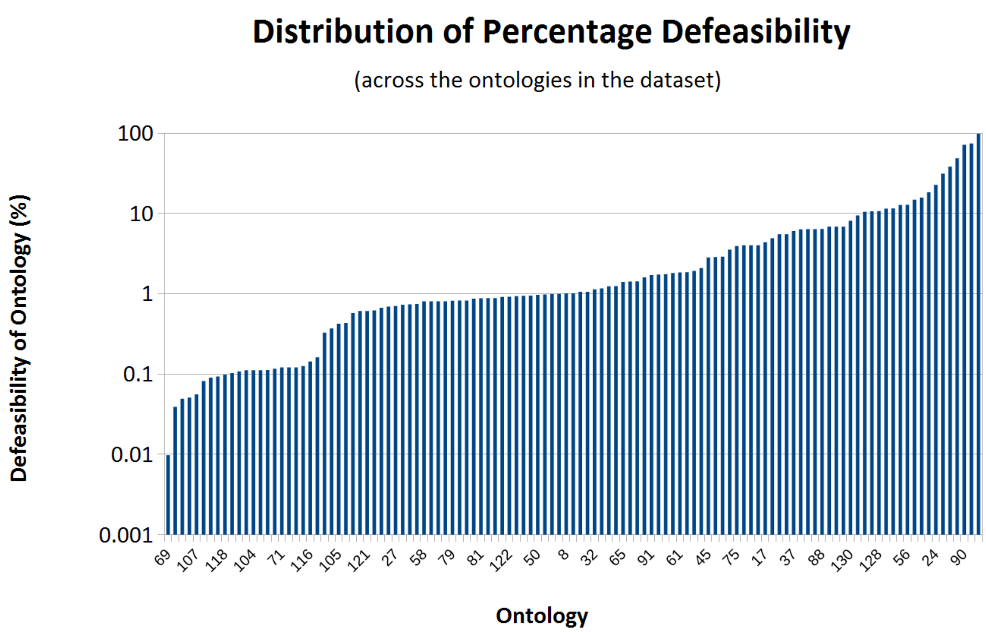


Figure 14: Percentage defeasibility distribution across the modified real world ontologies.

It is very interesting to note that the percentage defeasibility of most ontologies in the data stay well below 10%. It would be interesting to note if our reasoning performance for the 10% defeasibility category of the artificial data (Section 8.1) generalises to the current data as well. There a number of factors which make the current data different from the artificial data. The main one in terms of performance is probably ontology size. In our current data we have far larger ontologies in general than the artificially generated ones.

In terms of the $\mathcal{ALC}$ constructor distribution, Figure 15 shows that for the most part the numbers for the real world data closely match those of the artificial data (see Figure 4b). The only discrepancy is with universal role restrictions which occurr more frequently in the real world data than in our artificial data. While this is not ideal, we conjecture that overall this would not detract from the significance of the results for the artificial data.

| Negation (%) | | Conjunction (%) | | Disjunction (%) | | Existential (%) | | Universal (%) | |
|---|---|---|---|---|---|---|---|---|---|
| avg | median | avg | median | avg | median | avg | median | avg | median |
| 6.16 | 1.49 | 21.07 | 11.06 | 4.25 | 2.2 | 26.55 | 17.81 | 13.97 | 4.29 |

Figure 15: $\mathcal{ALC}$ constructor distribution per ontology in our modified real world dataset.

DL expressivity of the data ranges from variants of $\mathcal{ALC}$ all the way up to $\mathcal{SROIQ}$ (Horrocks, Kutz, & Sattler, 2006). There are 35 DL variants in total represented in the data.

Additionally, we generated a set of entailment queries (defeasible subsumptions of the form $C \mathrel{\sqsubset\!\!\!\sim} D$) for each ontology. We use the same methodology for generating these queries, as we used for generating them for the artificial data.

The tasks that we execute, the optimisations employed as well as the physical equipment used to perform the evaluation are the same as for the artificial data.

### 8.7 Ranking Compilation Results

It must be pointed out that the presentation of our results for this data is going to be significantly different to that of the artificial data in Section 8.1. Of course, one reason is that we do not have the ontologies binned into neat categories w.r.t. percentage defeasibility.

The overall results for ranking compilation are quite promising. Figures 16 and 17 give an overview of some of the more pertinent results w.r.t. ranking compilation.

| TBox size | | | DTBox size | | | LHS-incoherent classes | | | Totally exceptional LHS classes | | Ranking size | | Ranking Time (s) | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| avg | median | max | avg | median | max | avg | median | max | avg | max | avg | max | avg | median | max |
| 26,089 | 861 | 415,258 | 354 | 39 | 7,842 | 138 | 4.5 | 4,716 | 12 | 322 | 1.5 | 3 | 176.40 | 0.21 | 16,263 |

Figure 16: Ontology metrics and ranking results for the modified real world data.

Examining the ranking times in Figures 16, we notice that on average over the entire dataset, it takes 3 minutes to rank a single ontology. However, looking at the "median" column of the ranking times, shows us that the majority of rankings were computed in less than half a second. The most intensive ranking to compute was Ontology 134 which has $415,258$ TBox axioms of which $6,010$ are defeasible (it took roughly 4 hours to compute).
We have very little variance in ranking size (between 1 and 3), therefore we do not need to examine in detail the influence of ranking size on the compilation time.

However, the most challenging cases in theory for our reasoner are the ones in which there are hidden strict inclusions in the DTBox. Examining the number of recursions

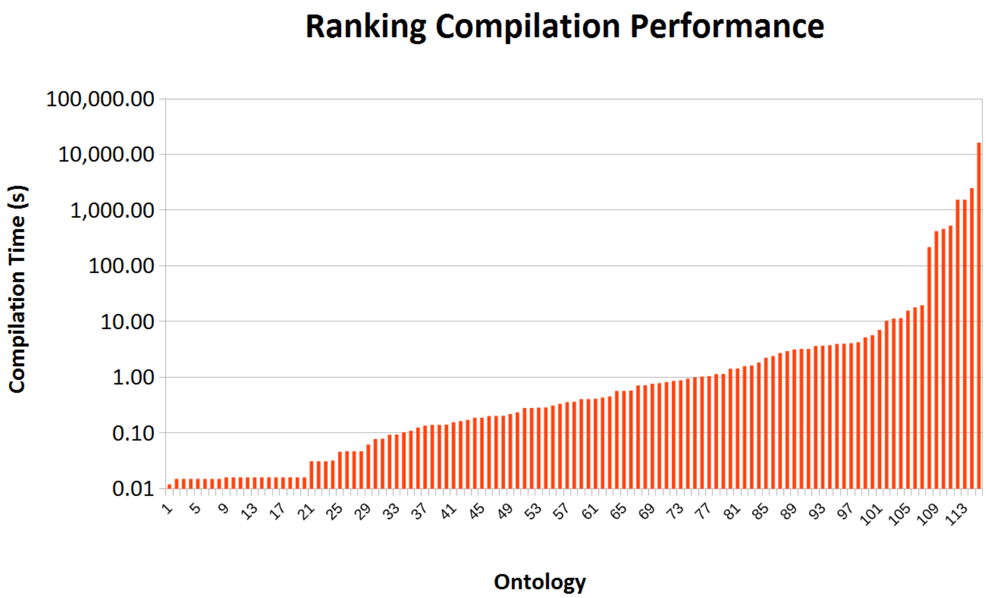


Figure 17: Ranking compilation time per modified real world ontology.

we have to perform over the data, we find that the need to recursively execute Procedure ComputeRanking is much less frequent than the artificial data (see Figure 18a).

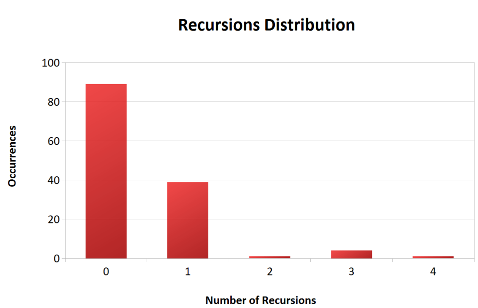


(a) Number of recursions required to rank the modified real world ontologies.

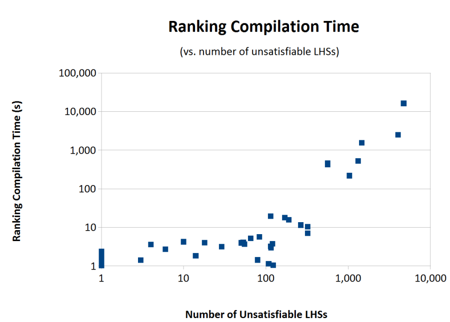


(b) Performance of ranking compilation vs. number of defeasible axioms in the ontology that have unsatisfiable LHSs.

Figure 18: Recursions and performance of ranking procedure on modified real world ontologies.

Therefore, we suspect that the number of recursions does not impact the hardness of ranking compilation considerably (with the current dataset). This is confirmed by the fact that the average compilation time for the cases with no recursions is by far the highest (255 seconds). Hardness, therefore, is mainly determined by other factors.

Since a naïve ranking compilation procedure has to check exceptionality of each defeasible axiom in the ontology, in most cases we expect the number of defeasible axioms to be the main contributor to hardness. However we also have an optimisation (Lemma 6) which says that we only need to check exceptionality of the defeasible axioms with unsatisfiable LHSs.

Therefore, it stands to reason that the number of unsatisfiable LHSs in the ontology would be the main contributor to hardness for our dataset. We plot the number of defeasible axioms in the ontology that have unsatisfiable LHSs against the ranking compilation time to reveal that this is indeed the case (see Figure 18b).

Both the X and Y axes are represented in logarithmic scale and we can see that from around 100 unsatisfiable LHSs the ranking times start to increase dramatically. In summary the compilation times for the modified real world data are, in general, comparable with those for our artificial data.

The average time to compile a ranking in the artificial data was around 100 seconds, whereas in our modified real world data this figure is around 176 seconds. However, we have much smaller percentage defeasibility ratios in the latter dataset than we do in the former. It would be interesting to see in the future whether real world defeasible ontologies would have similar percentage defeasibility ratios to those in our dataset.

An analysis of our algorithm, together with the results obtained in this evaluation, reveals that the number of unsatisfiable LHSs (and to a lesser extent the number of recursions) are the main contributors to hardness of ranking compilation.

It must also be mentioned that one should not ignore the number of strict axioms (i.e., overall ontology size with both defeasible and strict axioms) as a contributor to reasoning hardness. As we have repeatedly stressed, our algorithms are built upon classical entailment checks for which ontology size is the dominant indicator of hardness.

In fact, we notice that the average number of defeasible axioms in our real world dataset is only one third that of the artificial data, and yet we still obtained some cases in the real world data where compilation is harder than in the artificial data. We attribute most of this to the fact that ontology sizes are much larger on average in the real world dataset (see Figures 16 and 13 for a comparison). To conclude the ranking compilation analysis, we give some average metrics of this part of the evaluation in Figure 19.

| Defeasible Axioms | Ranking Size | Hidden Strict. | Rank 0 Size | Rank Size | Exception. Checks | Exceptional LHSs | Unsat. LHSs | Ranking Time (s) |
|---|---|---|---|---|---|---|---|---|
| 315.39 | 1.5 | 11.9 | 264.18 | 151.06 | 227.79 | 39.33 | 138.27 | 176.38 |

Figure 19: Average metrics obtained during the evaluation of ranking compilation performance for the modified real world data.

### 8.8 Entailment Checking Results

For Rational Closure, all queries (except those for Ontology 134) in the modified real world data, could be executed in less than a second. On average, over all ontologies, the query time was around 80ms and 90% of all queries could be executed in 200ms or less. A plot of the average query times for Rational Closure are presented in Figure 20.

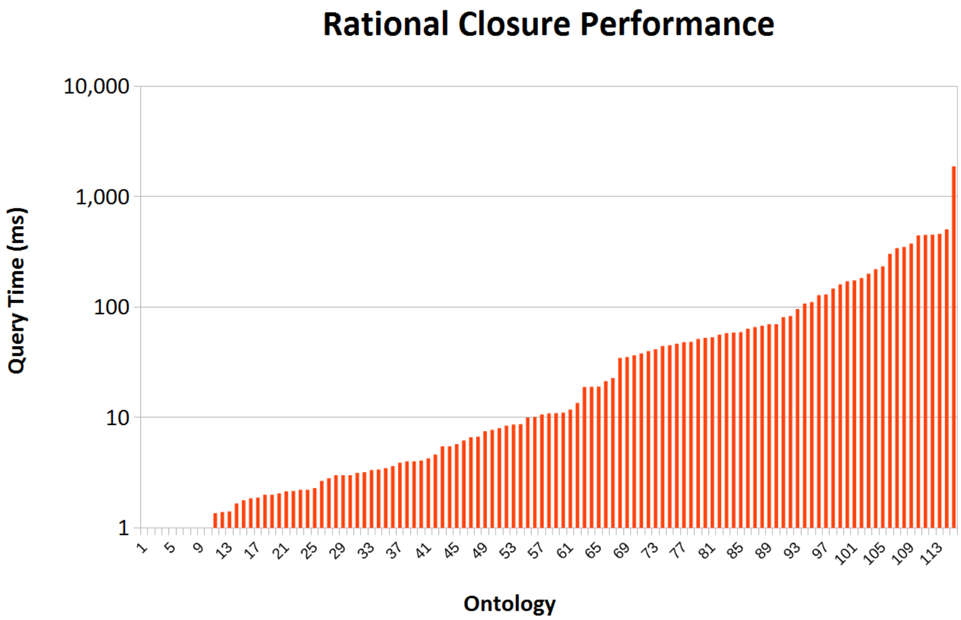


Figure 20: Rational closure performance in the modified real world data.

The data therefore confirms our analysis and generalisations in Section 8.1: that the performance of Rational closure (even using our preliminary implementation) is feasible for TBox reasoning in modern ontology editing tools. The vast majority of queries terminate within 100ms. There are, however, a significant number of queries which take between 100 and 500ms to compute. This is in slight contrast to the results for the artificial data which found that less than 1% of all queries posed to the reasoner took longer than 100ms to compute (extremely few queries even approached close to 100ms).

We hypothesise that the main reason for the queries between 100 and 500ms is the much larger ontology sizes obtained in the current dataset. In fact, we postulate that the sheer magnitude of ontology sizes has the largest impact on the performance of Rational Closure. Figure 21 lends credence to this claim (both axes are of logarithmic scale).

However, even though the performance of Rational Closure decreases considerably with the larger ontology sizes (we suspect the reason for this is that the performance of classical entailment also decreases as ontologies become larger), the performance still remains acceptable for practical reasoning tasks in ontology editing tools.

Just like in Section 8.1 we would like to ascertain, for the current dataset, how much more expensive the Rational Closure is than classical entailment. We find that, on average, the number of classical entailment checks required to check a defeasible entailment (using Rational closure) is 2.7. Since the most intensive component, by far, of the Rational closure algorithm is its classical entailment checks, we can infer that Rational closure takes roughly 2.7 times as long as classical entailment over the data. This is figure is quite similar to the one obtained in the artificial data - 3.5.

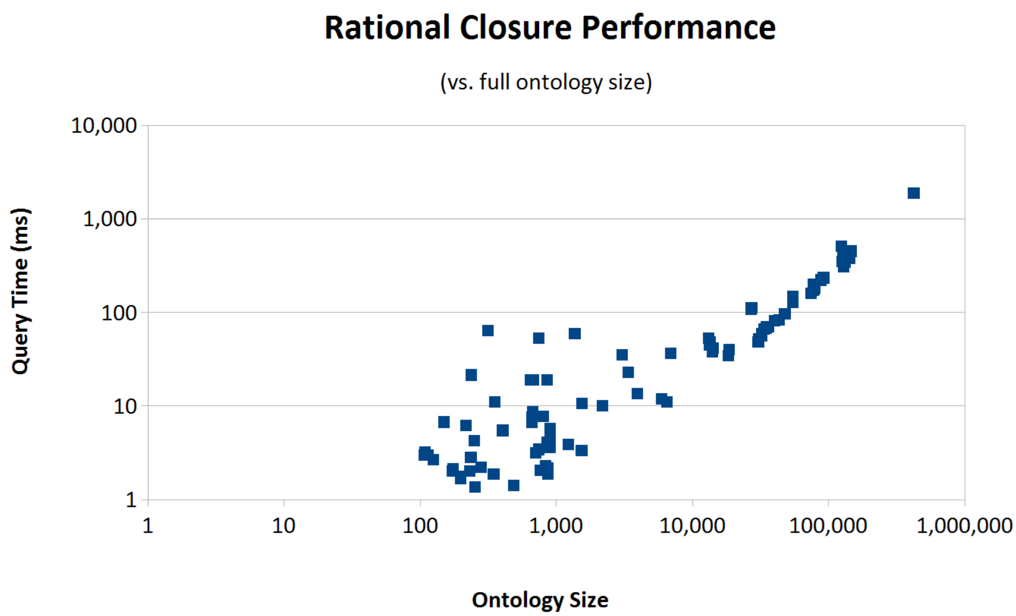


Figure 21: Average Rational closure performance vs. ontology size in the modified real world dataset.

## 9. Discussion and Related Work

Quantz and Royer (1992) were probably the first to consider the lifting of non-monotonic reasoning formalisms to a DL setting. They propose a general framework for Preferential Default Description Logics (PDDL) based on an $\mathcal{ALC}$-like language by introducing a version of default subsumption and proposing a semantics for it. Their semantics is based on a simplified version of standard DL interpretations in which all domains are assumed to be finite and the unique name assumption holds for object names. In that sense, their framework is much more restrictive than ours, as we do not make these assumptions here. They focus on a version of entailment which they refer to as preferential entailment, but which is to be distinguished from the version of preferential entailment that we have presented in this paper. In what follows, we shall refer to their version as *QR-preferential entailment*.

QR-preferential entailment is concerned with what ought to follow from a set of DL statements, together with a set of default subsumption statements, and is parameterised by a fixed partial order on (simplified) DL interpretations. They prove that any QR-preferential entailment satisfies the properties of a preferential consequence relation and, with some restrictions on the partial order, satisfies Rational Monotony as well. QR-preferential entailment can therefore be viewed as something in between the notions of preferential entailment (or rank entailment) and minimal rank entailment. It is also worth noting that although the QR-preferential entailments satisfy the properties of a preferential consequence relation, Quantz and Royer do not prove that QR-preferential entailment provides a characterisation of preferential consequence.

Other proposals to include default-style rules into description logics include the work of Baader and Hollunder (1995) and Padgham and Zhang (1993).

Closely related to our work is that of Giordano et al. (2009b) who use preferential orderings on $\Delta^{\mathcal{I}}$ to define a typicality operator $\mathbf{T}$ for $\mathcal{ALC}$ such that the expression $\mathbf{T}(C) \sqsubseteq D$ corresponds to our $C \precsim D$. They provide a version of a representation result for preferential orderings in terms of properties on selection functions (functions on the power set of the domain of interpretations), and present a tableaux calculus for computing preferential entailment that relies on KLM-style rules. Recently (Giordano et al., 2013b), they extended this work by considering modular orderings on $\Delta^{\mathcal{I}}$ (i.e., ranked interpretations) and then augment the inferential power of their system with a version of a minimal-model semantics, in which some ranked interpretations are preferred over others. This is similar in intuition to minimal rank entailment, but their approach also has a circumscriptive flavour to it (see below) since it relies on the specification of a set of concepts for which atypical instances must be minimized. As mentioned in Section 5, minimal rank entailment for $\mathcal{ALC}$ is based on the definition of minimal rank entailment for the propositional case, first presented by Giordano et al. (2012). In two recent papers (Giordano, Gliozzi, Olivetti, & Pozzato, 2013c; Giordano et al., 2013b) they extended this to the case for $\mathcal{ALC}$.

Casini and Straccia (2010, 2011) present KLM-based decision procedures for $\mathcal{ALC}$. Their proposal has a syntactic characterization, but lacks an appropriate semantics, a deficiency that this paper remedies. Lukasiewicz (2008) proposes probabilistic versions of the description logics $\mathcal{SHIF}(\mathbf{D})$ and $\mathcal{SHOIN}(\mathbf{D})$. As a special case of these logics, he obtains a version of a logic with defeasible subsumption with a semantics based on that of the propositional version of lexicographic closure (Lehmann, 1995).

Outside the family of preferential systems, there are mature proposals based on circumscription for DLs (Bonatti et al., 2009, 2011b, 2011a; Sengupta et al., 2011). The main drawback of these approaches is the burden on the ontology engineer to make appropriate decisions related to the (circumscriptive) fixing and varying of concepts and the priority of defeasible subsumption statements. Such choices can have a major effect on the conclusions drawn from the system, and can easily lead to counter-intuitive conclusions. Moreover, the use of circumscription usually implies a considerable increase in computational complexity w.r.t. the underlying monotonic entailment relation. The comparison between the present work and proposals outside the preferential family is more an issue about the pros and cons of the different kinds of non-monotonic reasoning, rather than about their DL re-formulation. As stated in the introduction, the preferential approach has a series of desirable qualities that, to our knowledge, no other approach to non-monotonic reasoning shares.

A more recent proposal is the novel approach proposed by Bonatti et al. (Bonatti et al., 2015a). Such a system, $DL^{\mathrm{N}}$, is not based on the preferential approach: its entailment relation does not satisfy the properties we have taken under consideration, but it is computationally tractable for any tractable classical DL.

Britz and Varzinczak (2012, 2013) explore the notion of *defeasible modalities*, with which defeasible effects of actions, defeasible knowledge, obligations and others can be formalized. Their approach differs from ours in that it is only preferential, but the semantic constructions are similar. This was recently extended (Britz, Casini, Meyer, & Varzinczak, 2013) to a notion of *defeasible role restrictions* in a DL setting. The idea comprises extending the language of $\mathcal{ALC}$ with an additional construct $\forall\!\!\!\!\sim$. The semantics of a concept $\tilde{\forall} r.C := \{x \in \Delta^{\mathcal{P}} \mid \min_{\prec_{\mathcal{P}}}(r^{\mathcal{P}}(x)) \subseteq C^{\mathcal{P}}\}$ is then given by all objects of $\Delta^{\mathcal{P}}$ such that all of their *minimal* $r$-related objects are $C$-instances. This is useful in situations where certain concept descriptions may be too strong. For example, the concept description $\mathsf{Lawyer} \sqcap \forall \mathsf{hasClient.PayingClient}$ would be too strong to capture the concept of lawyers who normally defend only paying clients, but who may exceptionally take on pro bono work, whereas the concept description $\mathsf{Lawyer} \sqcap \tilde{\forall} \mathsf{hasClient.PayingClient}$ would do the trick.

## 10. Conclusions and Future Work

The contributions of this paper are as follows: (*i*) The analysis of a simple and intuitive semantics for defeasible subsumption in description logics that is general enough to constitute the core framework within which to investigate non-monotonic extensions of DLs; (*ii*) A characterization of preferential and rational conditionals, with the respective representation results; (*iii*) An analysis of what an appropriate notion of entailment in a defeasible DL context mean and the analysis of a suitable candidate, namely minimal rank entailment, (*iv*) The formal connection between minimal rank entailment, the notion of rational closure and a syntactic method for its computation; (*v*) A thorough empirical evaluation of the cost of dealing with defeasible subsumption, both on synthetic ontologies and on existing ontologies.

An important consideration for future work is to extend the results obtained for $\mathcal{ALC}$ to more expressive logics. More specifically, the definition of minimal rank entailment would need to be revisited in many other DLs. The current definition, provided by Giordano et al. (2013c, 2013b) is based on *finite* ranked interpretations. The authors show that $\mathcal{ALC}$ with defeasible subsumption (strictly speaking, their logic $\mathcal{ALC}^{R}\mathbf{T}$) has the finite model property (if a DTBox has a ranked model, it has a finite ranked model). Because of this, rank entailment can be reduced to considering only finite ranked interpretations. For DLs without the finite model property, the definition of minimal rank entailment would have to be adapted in an appropriate manner.

Further topics for future research include the integration of notions such as *typicality* (Booth et al., 2012, 2013; Giordano et al., 2009b) and the aforementioned defeasible role restrictions into the framework here presented. Another avenue for future exploration is the study of belief revision for DLs via our results for rationality, mimicking the well-known link between belief revision and rational consequence in the propositional case (Gärdenfors & Makinson, 1994), thereby pushing the frontiers of defeasible reasoning in logics that are more expressive than the propositional one.

## References


Baader, F., Calvanese, D., McGuinness, D., Nardi, D., & Patel-Schneider, P. (Eds.). (2007). *The Description Logic Handbook: Theory, Implementation and Applications* (2 edition). Cambridge University Press.

Baader, F., & Hollunder, B. (1993). How to prefer more specific defaults in terminological default logic. In Bajcsy, R. (Ed.), *Proc. of the 13th Int. Joint Conf. on Artificial Intelligence (IJCAI'93)*, pp. 669–675. Morgan Kaufmann Publishers.

Baader, F., & Hollunder, B. (1995). Embedding defaults into terminological knowledge representation formalisms. *Journal of Automated Reasoning*, *14*(1), 149–180.

Baltag, A., & Smets, S. (2006). Dynamic belief revision over multi-agent plausibility models. In van der Hoek, W., & Wooldridge, M. (Eds.), *Logic and the Foundations of Game and Decision Theory (LOFT6)*, pp. 11–24. University of Liverpool.

Baltag, A., & Smets, S. (2008). A qualitative theory of dynamic interactive belief revision. In Bonanno, G., van der Hoek, W., & Wooldridge, M. (Eds.), *Logic and the Foundations of Game and Decision Theory (LOFT7)*, No. 3 in Texts in Logic and Games, pp. 13–60. Amsterdam University Press.

Bekiari, C., Doerr, M., Allocca, C., Barde, J., & Minadakis, N. (2014). *MarineTLO Documentation*. Forth.

Bonatti, P., Faella, M., & Sauro, L. (2011a). Defeasible inclusions in low-complexity DLs. *J. of Artificial Intelligence Research*, *42*, 719–764.

Bonatti, P., Faella, M., & Sauro, L. (2011b). On the complexity of EL with defeasible inclusions. In *Proc. of the 22nd Int. Joint Conf. on Artificial Intelligence (IJCAI 2011)*, pp. 762–767.

Bonatti, P., Lutz, C., & Wolter, F. (2009). The complexity of circumscription in description logic. *J. of Artificial Intelligence Research*, *35*, 717–773.

Bonatti, P. A., Faella, M., Petrova, I. M., & Sauro, L. (2015a). A new semantics for overriding in description logics. *Artificial Intelligence*, *222*, 1–48.

Bonatti, P. A., Petrova, I. M., & Sauro, L. (2015b). Optimizing the computation of overriding. In Arenas, M., Corcho, O., Simperl, E., Strohmaier, M., d'Aquin, M., Srinivas, K., Groth, P., Dumontier, M., Heflin, J., Thirunarayan, K., & Staab, S. (Eds.), *Proc. of the 14th International Semantic Web Conference (ISWC 2015)*, pp. 356–372, Cham. Springer.

Booth, R., Meyer, T., & Varzinczak, I. (2012). PTL: A propositional typicality logic. In Fariñas del Cerro, L., Herzig, A., & Mengin, J. (Eds.), *Proc. of the 13th Eur. Workshop on Logics in Artificial Intelligence (JELIA'2012)*, No. 7519 in LNCS, pp. 107–119. Springer.

Booth, R., Meyer, T., & Varzinczak, I. (2013). A propositional typicality logic for extending rational consequence. In Fermé, E., Gabbay, D., & Simari, G. (Eds.), *Trends in Belief Revision and Argumentation Dynamics*, Vol. 48 of *Studies in Logic – Logic and Cognitive Systems*, pp. 123–154. King's College Publications.

Booth, R., & Paris, J. (1998). A note on the rational closure of knowledge bases with both positive and negative knowledge. *Journal of Logic, Language and Information*, *7*(2), 165–190.

Boutilier, C. (1994). Conditional logics of normality: A modal approach. *Artificial Intelligence*, *68*(1), 87–154.

Britz, K., Casini, G., Meyer, T., & Varzinczak, I. (2013). Preferential role restrictions. In *Proc. of the 2013 Description Logic Workshop (DL 2013)*, pp. 93–106.

Britz, K., Heidema, J., & Meyer, T. (2008). Semantic preferential subsumption. In Lang, J., & Brewka, G. (Eds.), *Proc. of the 11th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR 2008)*, pp. 476–484. AAAI Press/MIT Press.

Britz, K., Heidema, J., & Meyer, T. (2009). Modelling object typicality in description logics. In Nicholson, A., & Li, X. (Eds.), *Proceedings of the 22nd Australasian Joint Conference on Artificial Intelligence*, No. 5866 in LNAI, pp. 506–516. Springer.

Britz, K., Meyer, T., & Varzinczak, I. (2011a). Preferential reasoning for modal logics. *Electronic Notes in Theoretical Computer Science*, *278*, 55–69. Proceedings of the 7th Workshop on Methods for Modalities (M4M'2011).

Britz, K., Meyer, T., & Varzinczak, I. (2011b). Semantic foundation for preferential description logics. In Wang, D., & Reynolds, M. (Eds.), *Proceedings of AI*, No. 7106 in LNAI, pp. 491–500. Springer.

Britz, K., Meyer, T., & Varzinczak, I. (2012). Normal modal preferential consequence. In Thielscher, M., & Zhang, D. (Eds.), *Proceedings of the 25th Australasian Joint Conference on Artificial Intelligence*, No. 7691 in LNAI, pp. 505–516. Springer.

Britz, K., & Varzinczak, I. (2012). Defeasible modes of inference: A preferential perspective. In *Proceedings of the 14th International Workshop on Nonmonotonic Reasoning (NMR)*.

Britz, K., & Varzinczak, I. (2013). Defeasible modalities. In *Proceedings of the 14th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, pp. 49–60.

Casini, G., Meyer, T., Moodley, K., & Varzinczak, I. (2013a). Nonmonotonic reasoning in description logics: Rational closure for the Abox. In *Proc. of the 2013 Description Logic Workshop (DL 2013)*, pp. 600–615.

Casini, G., Meyer, T., Moodley, K., & Varzinczak, I. (2013b). Towards practical defeasible reasoning for description logics. In *Proc. of the 2013 Description Logic Workshop (DL 2013)*, pp. 587–599.

Casini, G., & Straccia, U. (2010). Rational closure for defeasible description logics. In Janhunen, T., & Niemelä, I. (Eds.), *Proc. of the 12th Eur. Workshop on Logics in Artificial Intelligence (JELIA'2010)*, No. 6341 in LNCS, pp. 77–90. Springer-Verlag.

Casini, G., & Straccia, U. (2014). Lexicographic closure for defeasible Description Logics. In *Proceedings of the 8th Australasian Ontology Workshop (AOW)*, Vol. 969, pp. 4–15. CEUR Workshop Proceedings.

Casini, G., Meyer, T., Moodley, K., Sattler, U., & Varzinczak, I. (2015). Introducing Defeasibility into OWL Ontologies. In *Proc. of the 14th International Semantic Web Conference (ISWC 2015)*, pp. 409–426. Springer.

Chellas, B. (1980). *Modal logic: An introduction.* Cambridge University Press.

Donini, F., & Massacci, F. (2000). EXPTIME tableaux for $\mathcal{ALC}$. *Artificial Intelligence*, *124*, 87–138.

Donini, F., Nardi, D., & Rosati, R. (2002). Description logics of minimal knowledge and negation as failure. *ACM Transactions on Computational Logic*, *3*(2), 177–225.

Gärdenfors, P., & Makinson, D. (1994). Nonmonotonic inference based on expectations. *Artificial Intelligence*, *65*(2), 197–245.

Giordano, L., Gliozzi, V., Olivetti, N., & Pozzato, G. (2007). Preferential description logics. In Dershowitz, N., & Voronkov, A. (Eds.), *Proc. of the 14th Int. Conf. on Logic for Programming, Artificial Intelligence and Reasoning (LPAR 2007)*, No. 4790 in LNAI, pp. 257–272. Springer.

Giordano, L., Gliozzi, V., Olivetti, N., & Pozzato, G. (2008). Reasoning about typicality in preferential description logics. In Hölldobler, S., Lutz, C., & Wansing, H. (Eds.), *Proc. of the 11th Eur. Workshop on Logics in Artificial Intelligence (JELIA'2008)*, No. 5293 in LNAI, pp. 192–205. Springer.

Giordano, L., Gliozzi, V., Olivetti, N., & Pozzato, G. (2009a). Analytic tableaux calculi for KLM logics of nonmonotonic reasoning. *ACM Transactions on Computational Logic*, *10*(3), 18:1–18:47.

Giordano, L., Gliozzi, V., Olivetti, N., & Pozzato, G. (2009b). $\mathcal{ALC} + T$: a preferential extension of description logics. *Fundamenta Informaticae*, *96*(3), 341–372.

Giordano, L., Gliozzi, V., Olivetti, N., & Pozzato, G. (2012). A minimal model semantics for nonmonotonic reasoning. In Fariñas del Cerro, L., Herzig, A., & Mengin, J. (Eds.), *Proceedings of JELIA*, No. 7519 in LNCS, pp. 228–241. Springer.

Giordano, L., Gliozzi, V., Olivetti, N., & Pozzato, G. (2013a). Minimal model semantics and rational closure in description logics. In *Proc. of the 2013 Description Logic Workshop (DL 2013)*.

Giordano, L., Gliozzi, V., Olivetti, N., & Pozzato, G. (2013b). A non-monotonic description logic for reasoning about typicality. *Artificial Intelligence*, *195*, 165–202.

Giordano, L., Gliozzi, V., Olivetti, N., & Pozzato, G. (2013c). Rational closure in description logics of typicality. In *Proceedings of IAF*.

Giordano, L., Gliozzi, V., Olivetti, N., & Pozzato, G. (2015). Semantic characterization of rational closure: From propositional logic to description logics. *Artificial Intelligence*, *226*, 1–33.

Glimm, B., Horrocks, I., Motik, B., Stoilos, G., & Wang, Z. (2014). HermiT: an OWL 2 reasoner. *Journal of Automated Reasoning*, *53*(3), 245–269.

Grau, B. C., Horrocks, I., Kazakov, Y., & Sattler, U. (2008). Modular reuse of ontologies: Theory and practice.. *J. of Artificial Intelligence Research*, *31*, 273–318.

Guo, Y., Pan, Z., & Heflin, J. (2005). LUBM: A Benchmark for OWL Knowledge Base Systems. *Web Semantics: SSAWWW*, *3*(2), 158–182.

Haarslev, V., & Möller, R. (2001). High Performance Reasoning with Very Large Knowledge Bases: A Practical Case Study. In *Proc. of the 17th Int. Joint Conf. on Artificial Intelligence (IJCAI 2001)*, Vol. 1, pp. 161–168.

Hoehndorf, R., Loebe, F., Kelso, J., & Herre, H. (2007). Representing Default Knowledge in Biomedical Ontologies: Application to the Integration of Anatomy and Phenotype Ontologies. *BMC Bioinformatics*, *8*(1), 377.

Horridge, M., & Bechhofer, S. (2011). The OWL API: A Java API for OWL Ontologies. *Semantic Web*, *2*(1), 11–21.

Horridge, M. (2011). *Justification based explanation in ontologies.* Ph.D. thesis, The University of Manchester.

Horrocks, I., Kutz, O., & Sattler, U. (2006). The Even More Irresistible $\mathcal{SROIQ}$. In *Proc. of the 10th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR 2006)*, pp. 57–67.

Kraus, S., Lehmann, D., & Magidor, M. (1990). Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, *44*, 167–207.

Lehmann, D. (1989). What does a conditional knowledge base entail?. In Brachman, R., & Levesque, H. (Eds.), *Proc. of the 1st Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR'89)*, pp. 212–222.

Lehmann, D. (1995). Another perspective on default reasoning. *Annals of Mathematics and Artificial Intelligence*, *15*(1), 61–82.

Lehmann, D., & Magidor, M. (1992). What does a conditional knowledge base entail?. *Artificial Intelligence*, *55*, 1–60.

Lukasiewicz, T. (2008). Expressive probabilistic description logics. *Artificial Intelligence*, *172*(6-7), 852–883.

Ma, L., Yang, Y., Qiu, Z., Xie, G., Pan, Y., & Liu, S. (2006). Towards a Complete OWL Ontology Benchmark. In *The Semantic Web: Research and Applications*, pp. 125–139. Springer Berlin.

Matentzoglu, N., Bail, S., & Parsia, B. (2013). A Snapshot of the OWL Web. In *Proc. of the 2013 International Semantic Web Conference (ISWC 2013)*, pp. 331–346.

Matentzoglu, N., Tang, D., Parsia, B., & Sattler, U. (2014a). The Manchester OWL Repository: System Description. In *Proc. of the 2014 International Semantic Web Conference (ISWC 20014)*, pp. 285–288.

Matentzoglu, N., Tang, D., Parsia, B., & Sattler, U. (2014b). The Manchester OWL Repository: System Description. In *Proc. of the 2014 International Semantic Web Conference (ISWC 20014)*, pp. 285–288.

McGuinness, D. L., Van Harmelen, F., et al. (2004). OWL Web Ontology Language Overview. *W3C recommendation*, *10*(10).

Padgham, L., & Zhang, T. (1993). A terminological logic with defaults: A definition and an application. In Bajcsy, R. (Ed.), *Proc. of the 13th Int. Joint Conf. on Artificial Intelligence (IJCAI'93)*, pp. 662–668. Morgan Kaufmann Publishers.

Quantz, J., & Royer, V. (1992). A preference semantics for defaults in terminological logics. In *Proc. of the 3rd Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR'92)*, pp. 294–305.

Quantz, J., & Ryan, M. (1993). Preferential default description logics. Tech. rep., TU Berlin.

Raskin, R. G., & Pan, M. J. (2005). Knowledge Representation in the Semantic Web for Earth and Environmental Terminology (SWEET). *Computers & Geosciences*, *31*(9), 1119–1125.

Rector, A. L. (2004). Defaults, Context, and Knowledge: Alternatives for OWL-indexed Knowledge Bases. In *Pacific Symposium on Biocomputing*, Vol. 9, pp. 226–237. World Scientific.

Romero, A. A., Cuenca-Grau, B., & Horrocks, I. (2012). MORe: Modular Combination of OWL Reasoners for Ontology Classification. In *Proc. of the 2012 International Semantic Web Conference (ISWC 20012)*, pp. 1–16.

Rott, H. (2001). *Change, Choice and Inference: a study of belief revision and nonmonotonic reasoning.* Oxford University Press, Oxford, UK.

Sattler, U., Schneider, T., & Zakharyaschev, M. (2009). Which Kind of Module Should I Extract?. In *Proc. of the 2009 Description Logic Workshop (DL 2009).*

Sazonau, V., Sattler, U., & Brown, G. (2014). Predicting OWL Reasoners: Locally or Globally?. In *Proc. of the 2014 Description Logic Workshop (DL 2014)*, pp. 713–724.

Schild, K. (1991). A correspondence theory for terminological logics: Preliminary report. In *Proc. of the 12th Int. Joint Conf. on Artificial Intelligence (IJCAI'91)*, pp. 466–471.

Schulz, S., Stenzhorn, H., Boeker, M., & Smith, B. (2009). Strengths and Limitations of Formal Ontologies in the Biomedical Domain. *Revista Electronica de Comunicacao, Informacao & Inovacao em Saude: RECIIS*, *3*(1), 31.

Sengupta, K., Alfa Krisnadhi, A., & Hitzler, P. (2011). Local closed world semantics: Grounded circumscription for OWL. In Aroyo, L., Welty, C., Alani, H., Taylor, J., Bernstein, A., Kagal, L., Noy, N., & Blomqvist, E. (Eds.), *Proc. of the 2011 International Semantic Web Conference (ISWC 2011)*, pp. 617–632.

Shoham, Y. (1988). *Reasoning about Change: Time and Causation from the Standpoint of Artificial Intelligence.* MIT Press.

Sioutos, N., Coronado, S. d., Haber, M. W., Hartel, F. W., Shaiu, W.-L., & Wright, L. W. (2007). NCI Thesaurus: A Semantic Model Integrating Cancer-related Clinical and Molecular Information. *Journal of Biomedical Informatics*, *40*(1), 30–43.

Steigmiller, A., Liebig, T., & Glimm, B. (2014). Konclude: System Description. *Web Semantics: Web Semantics: Science, Services and Agents on the World Wide Web*, *27*, 78–85.

Straccia, U., & Casini, G. (2011). Defeasible inheritance-based description logics. In *Proc. of the 22nd Int. Joint Conf. on Artificial Intelligence (IJCAI 2011).*

Thomas, E., Pan, J. Z., & Ren, Y. (2010). TrOWL: Tractable OWL 2 Reasoning Infrastructure. In *Proc. of the 2010 International Semantic Web Conference (ISWC 2010)*, pp. 431–435. Springer.

Tsarkov, D., & Horrocks, I. (2006). FaCT++ Description Logic Reasoner: System description. In *Automated Reasoning*, pp. 292–297. Springer.

Tsarkov, D., & Palmisano, I. (2012). Chainsaw: a Metareasoner for Large Ontologies. In *Proceedings of the OWL Reasoner Evaluation Workshop.*

Tzitzikas, Y., Allocca, C., Bekiari, C., Marketakis, Y., Fafalios, P., Doerr, M., Minadakis, N., Patkos, T., & Candela, L. (2013). Integrating Heterogeneous and Distributed

Information about Marine Species through a Top Level Ontology. In Garoufallou, E., & Greenberg, J. (Eds.), *Metadata and Semantics Research*, Vol. 390 of *Communications in Computer and Information Science*, pp. 289–301. Springer International Publishing.

Wandelt, S., & Möller, R. (2012). Towards abox modularization of semi-expressive description logics. *Applied Ontology*, *7*(2), 133–167.

Whetzel, P. L., Noy, N. F., Shah, N. H., Alexander, P. R., Nyulas, C., Tudorache, T., & Musen, M. A. (2011). Bioportal: enhanced functionality via new web services from the national center for biomedical ontology to access and use ontologies in software applications. *Nucleic acids research*, *39*(suppl 2), W541–W545.

## Appendix A. Proofs for Section 3

**Note:** the propositions associated with the symbol $(*)$ are introduced here in the appendix, while they are omitted in the main text.

### A.1 Proof of Theorem 1

SOUNDNESS

(Cons): Assume that $\top \rightsquigarrow_{\mathcal{P}} \bot$ is the case. Then $\mathcal{P} \Vdash \top \sqsubsetsim \bot$, i.e., $\min_{\prec} \top^{\mathcal{P}} \subseteq \bot^{\mathcal{P}}$, and therefore $\min_{\prec} \mathcal{D}^{\mathcal{P}} \subseteq \emptyset$. This implies $\min_{\prec} \mathcal{D}^{\mathcal{P}} = \emptyset$, and then $\mathcal{D}^{\mathcal{P}} = \emptyset$, which is a contradiction. Hence $\top \not\rightsquigarrow_{\mathcal{P}} \bot$.

(Ref): Let $x \in \mathcal{D}^{\mathcal{P}}$ be such that $x \in \min_{\prec} C^{\mathcal{P}}$. Then clearly $x \in C^{\mathcal{P}}$ and therefore $\mathcal{P} \Vdash C \sqsubsetsim C$. Hence $C \rightsquigarrow_{\mathcal{P}} C$.

(LLE): Assume that $C \rightsquigarrow_{\mathcal{P}} E$ and $\models C \equiv D$. Then $\mathcal{P} \Vdash C \sqsubsetsim E$, which means $\min_{\prec} C^{\mathcal{P}} \subseteq E^{\mathcal{P}}$. Since $\models C \equiv D$, in particular we have $\mathcal{P} \Vdash C \equiv D$, i.e., $C^{\mathcal{P}} = D^{\mathcal{P}}$, and therefore $\min_{\prec} C^{\mathcal{P}} = \min_{\prec} D^{\mathcal{P}}$. Hence $\min_{\prec} D^{\mathcal{P}} \subseteq E^{\mathcal{P}}$, and therefore $\mathcal{P} \Vdash D \sqsubsetsim E$, from which follows $D \rightsquigarrow_{\mathcal{P}} E$.

(And): Assume we have both $C \rightsquigarrow_{\mathcal{P}} D$ and $C \rightsquigarrow_{\mathcal{P}} E$. Then $\mathcal{P} \Vdash C \sqsubsetsim D$ and $\mathcal{P} \Vdash C \sqsubsetsim E$, i.e., $\min_{\prec} C^{\mathcal{P}} \subseteq D^{\mathcal{P}}$ and $\min_{\prec} C^{\mathcal{P}} \subseteq E^{\mathcal{P}}$, and then $\min_{\prec} C^{\mathcal{P}} \subseteq D^{\mathcal{P}} \cap E^{\mathcal{P}}$, from which follows $\min_{\prec} C^{\mathcal{P}} \subseteq (D \sqcap E)^{\mathcal{P}}$. Hence $\mathcal{P} \Vdash C \sqsubsetsim D \sqcap E$, and therefore $C \rightsquigarrow_{\mathcal{P}} D \sqcap E$.

(Or): Assume we have both $C \rightsquigarrow_{\mathcal{P}} E$ and $D \rightsquigarrow_{\mathcal{P}} E$. Let $x \in \min_{\prec}(C \sqcup D)^{\mathcal{P}}$. Then $x$ is minimal in $C^{\mathcal{P}} \cup D^{\mathcal{P}}$, and therefore either $x \in \min_{\prec} C^{\mathcal{P}}$ or $x \in \min_{\prec} D^{\mathcal{P}}$. In either case $x \in E^{\mathcal{P}}$. Hence $\mathcal{P} \Vdash C \sqcup D \sqsubsetsim E$ and therefore $C \sqcup D \rightsquigarrow_{\mathcal{P}} E$.

(RW): Assume we have both $C \rightsquigarrow_{\mathcal{P}} D$ and $\models D \sqsubseteq E$. Then $\mathcal{P} \Vdash C \sqsubsetsim D$ and $\mathcal{P} \Vdash D \sqsubseteq E$, i.e., $\min_{\prec} C^{\mathcal{P}} \subseteq D^{\mathcal{P}}$ and $D^{\mathcal{P}} \subseteq E^{\mathcal{P}}$. Hence $\min_{\prec} C^{\mathcal{P}} \subseteq E^{\mathcal{P}}$ and then $\mathcal{P} \Vdash C \sqsubsetsim E$. Therefore $C \rightsquigarrow_{\mathcal{P}} E$.

(CM): Assume we have both $C \rightsquigarrow_{\mathcal{P}} D$ and $C \rightsquigarrow_{\mathcal{P}} E$. Then $\mathcal{P} \Vdash C \sqsubsetsim D$ and $\mathcal{P} \Vdash C \sqsubsetsim E$, and therefore $\min_{\prec} C^{\mathcal{P}} \subseteq D^{\mathcal{P}}$ and $\min_{\prec} C^{\mathcal{P}} \subseteq E^{\mathcal{P}}$. Let $x \in \min_{\prec}(C \sqcap D)^{\mathcal{P}}$. We show that $x \in \min_{\prec} C^{\mathcal{P}}$. Suppose this is not the case. Since $\prec$ is smooth, there must be $x' \in \min_{\prec} C^{\mathcal{P}}$ such that $x' \prec x$. Because $\mathcal{P} \Vdash C \sqsubsetsim D$, $x' \in D^{\mathcal{P}}$, and then $x' \in C^{\mathcal{P}} \cap D^{\mathcal{P}}$, i.e., $x' \in (C \sqcap D)^{\mathcal{P}}$. From this and $x' \prec x$ it follows that $x$ is not minimal in $(C \sqcap D)^{\mathcal{P}}$, which is a contradiction. Hence $x \in \min_{\prec} C^{\mathcal{P}}$. From this and $\min_{\prec} C^{\mathcal{P}} \subseteq E^{\mathcal{P}}$, it follows that $x \in E^{\mathcal{P}}$. Hence $\mathcal{P} \Vdash C \sqcap D \sqsubsetsim E$, and therefore $C \sqcap D \rightsquigarrow_{\mathcal{P}} E$. ■

Completeness

Let $\leadsto \subseteq \mathcal{L} \times \mathcal{L}$ be a preferential conditional. We shall construct a preferential interpretation $\mathcal{P}$ such that $\leadsto_{\mathcal{P}} := \{(C, D) \mid \mathcal{P} \Vdash C \sqsubsetapprox D\} = \leadsto$.

**Definition 24** Let $\mathscr{U} := \{(\mathcal{I}, x) \mid \mathcal{I} = \langle \mathcal{D}^{\mathcal{I}}, \cdot^{\mathcal{I}} \rangle$ and $x \in \mathcal{D}^{\mathcal{I}}\}$.

Intuitively, $\mathscr{U}$ denotes the universe of objects in the context of their respective DL interpretations, i.e., $\mathscr{U}$ is a set of first-order interpretations.

**Definition 25** A pair $(\mathcal{I}, x) \in \mathscr{U}$ is normal for $C \in \mathcal{L}$ if for every $D \in \mathcal{L}$ such that $C \leadsto D$, $x \in D^{\mathcal{I}}$.

**Lemma 8 (∗)** *Let $\leadsto \subseteq \mathcal{L} \times \mathcal{L}$ satisfy (Ref), (RW) and (And), and let $C, D \in \mathcal{L}$. Then all normal $(\mathcal{I}, x)$ for $C$ satisfy $D$ if and only if $C \leadsto D$.*

**Proof:**
The 'if'-part follows from the definition of normality above. For the 'only if'-part, assume $C \not\leadsto D$. We build a pair $(\mathcal{I}, x)$ that is normal for $C$ but that does not satisfy $D$. Let $\Gamma := \{\neg D\} \cup \{E \mid C \leadsto E\}$. All we need to do is show that there is $(\mathcal{I}, x)$ such that $x \in F^{\mathcal{I}}$ for every $F \in \Gamma$. Suppose this is not the case. Then by compactness there exists a finite $\Gamma' \subseteq \Gamma$ such that $\models \bigsqcap_{F \in \Gamma'} F \sqsubseteq D$. From this follows $\models \top \sqsubseteq \neg \bigsqcap_{F \in \Gamma'} \sqcup D$, and, in particular, we have $\models C \sqsubseteq \neg \bigsqcap_{F \in \Gamma'} F \sqcup D$. Now from (Ref) we have $C \leadsto C$. From this, $\models C \sqsubseteq \neg \bigsqcap_{F \in \Gamma'} F \sqcup D$ and (RW) we get $C \leadsto (\neg \bigsqcap_{F \in \Gamma'} F \sqcup D)$. But we also have $C \leadsto \bigsqcap_{F \in \Gamma'} F$ by the (And) rule, and then by applying (And) once more we derive $C \leadsto \bigsqcap_{F \in \Gamma'} F \sqcap (\neg \bigsqcap_{F \in \Gamma'} F \sqcup D)$. From this and (RW) we conclude $C \leadsto D$, from which we derive a contradiction. ■

**Lemma 9 (∗)** *If $\leadsto$ is preferential, the following rule holds:*

$$\frac{C \sqcup D \leadsto C,\ \ D \sqcup E \leadsto D}{C \sqcup E \leadsto C}$$

**Proof:**
The proof is analogous to that by Kraus et al. (Kraus et al., 1990, Lemma 5.5) and we do not repeat it here. ■

**Definition 26** Let $C, D \in \mathcal{L}$. $C \leq D$ if $C \sqcup D \leadsto C$.

**Lemma 10 (∗)** *If $\leadsto$ is preferential, then $\leq$ is reflexive and transitive.*

**Proof:**
From (Ref) we have $C \leadsto C$. This and (LLE) gives us $C \sqcup C \leadsto C$, therefore we have $C \leq C$ and $\leq$ is reflexive. Transitivity follows from Lemma 9. ■

**Lemma 11 (∗)** *If $\leadsto$ is preferential, the following rule holds:*

$$\frac{C \sqcup D \leadsto C,\ \ D \leadsto E}{C \leadsto \neg D \sqcup E}$$

**Proof:**
The proof is analogous to that by Kraus et al. (Kraus et al., 1990, Lemma 5.5) and we do not repeat it here. ■

**Lemma 12** **($*$)** *If $C \leq D$ and $(\mathcal{I}, x)$ is normal for $C$, and $x \in D^{\mathcal{I}}$, then $(\mathcal{I}, x)$ is normal for $D$.*

**Proof:**
From $C \leq D$ we get $C \sqcup D \rightsquigarrow C$. Assume that $D \rightsquigarrow E$ is the case. Then by Lemma 11 we have $C \rightsquigarrow \neg D \sqcup E$. Since $(\mathcal{I}, x)$ is normal for $C$, we have $x \in (\neg D \sqcup E)^{\mathcal{I}}$. Given that $x \in D^{\mathcal{I}}$, we must have $x \in E^{\mathcal{I}}$. ■

**Lemma 13** **($*$)** *If $\rightsquigarrow$ is preferential, the following rule holds:*

$$\frac{C \sqcup D \rightsquigarrow C,\ D \sqcup E \rightsquigarrow D}{C \rightsquigarrow \neg E \sqcup D}$$

**Proof:**
The proof is analogous to that by Kraus et al. (Kraus et al., 1990, Lemma 5.5) and we do not repeat it here. ■

**Lemma 14** **($*$)** *If $C \leq D \leq E$ and $(\mathcal{I}, x)$ is normal for $C$, and $x \in E^{\mathcal{I}}$, then $(\mathcal{I}, x)$ is normal for $D$.*

**Proof:**
By Lemma 12, it is enough to show that $x \in D^{\mathcal{I}}$. By Lemma 13 we have $C \rightsquigarrow \neg E \sqcup D$. Since $(\mathcal{I}, x)$ is normal for $C$ and $x \in E^{\mathcal{I}}$, then we must have $x \in D^{\mathcal{I}}$. ■

We now construct a preferential interpretation as in Definition 1.

Let $\mathcal{C} := \{C \mid C \rightsquigarrow \bot\}$ and let $\mathscr{I} := \{\mathcal{I} = \langle \mathcal{D}^{\mathcal{I}}, \cdot^{\mathcal{I}} \rangle \mid C^{\mathcal{I}} = \emptyset$ for all $C \in \mathcal{C}\}$. Intuitively, $\mathscr{I}$ contains all interpretations that are 'compatible' with $\rightsquigarrow$ in the sense of not satisfying concepts that are defeasibly subsumed by the contradiction.

For each $\mathcal{I} \in \mathscr{I}$, let $\mathcal{I}^{+} = \langle \mathcal{D}^{\mathcal{I}^{+}}, \cdot^{\mathcal{I}^{+}} \rangle$ be such that:

- $\mathcal{D}^{\mathcal{I}^{+}} := X^{C} \cup X^{\bot}$, where $X^{C} := \{\langle \mathcal{I}, x, C \rangle \mid (\mathcal{I}, x)$ is normal for $C \in \mathcal{L}\}$, and $X^{\bot} := \{\langle \mathcal{I}, x, \bot \rangle \mid (\mathcal{I}, x)$ is not normal for any $C \in \mathcal{L}\}$;
- $\cdot^{\mathcal{I}^{+}}$ is such that for every $D \in \mathcal{L}$, $\langle \mathcal{I}, x, C \rangle \in D^{\mathcal{I}^{+}}$ if and only if $x \in D^{\mathcal{I}}$, and for every $r \in \mathsf{N}_{\mathscr{R}}$ $(\langle \mathcal{I}, x, C \rangle, \langle \mathcal{I}, y, D \rangle) \in r^{\mathcal{I}^{+}}$ if and only if $(x, y) \in r^{\mathcal{I}}$.

Let $\mathcal{P} := \langle \mathcal{D}^{\mathcal{P}}, \cdot^{\mathcal{P}}, \prec \rangle$ be such that:

- $\mathcal{D}^{\mathcal{P}} := \bigcup_{\mathcal{I} \in \mathscr{I}} \mathcal{D}^{\mathcal{I}^{+}}$;
- $\cdot^{\mathcal{P}} := \bigcup_{\mathcal{I} \in \mathscr{I}} \cdot^{\mathcal{I}^{+}}$;
- $\prec$ is the smallest relation such that:
  - For every $\langle \mathcal{I}, x, C \rangle \in \mathcal{D}^{\mathcal{P}}$ such that $C \neq \bot$, $\langle \mathcal{I}, x, C \rangle \prec \langle \mathcal{J}, y, \bot \rangle$ for every $\langle \mathcal{J}, y, \bot \rangle \in \mathcal{D}^{\mathcal{P}}$;

– For every $\langle \mathcal{I}, x, C\rangle, \langle \mathcal{J}, y, D\rangle \in \mathcal{D}^{\mathcal{P}}$ such that $C, D \neq \bot$, $\langle \mathcal{I}, x, C\rangle \prec \langle \mathcal{J}, y, D\rangle$ if and only if $C \leq D$ and $x \notin D^{\mathcal{I}}$.

(In the construction of $\mathcal{P}$, note that all pairs $(\mathcal{I}, x)$ that are not normal for any concept $C$ are moved higher up in the ordering so that they correspond to the least preferred objects of the domain.)

In Lemmas 15–20 below we show that $\mathcal{P}$ as constructed above is indeed a preferential interpretation.

**Lemma 15 (∗)** $\mathcal{D}^{\mathcal{P}} \neq \emptyset$.

**Proof:**
From $\top \not\rightsquigarrow \bot$ and Lemma 8, it follows that there is some normal $(\mathcal{I}, x)$ for $\top$ that does not satisfy $\bot$. Hence $\langle \mathcal{I}, x, \top\rangle \in \mathcal{D}^{\mathcal{P}}$ and therefore $\mathcal{D}^{\mathcal{P}} \neq \emptyset$. ■

**Lemma 16 (∗)** $C \leq \bot$ *for every* $C \in \mathcal{L}$.

**Proof:**
By (Ref) we have $C \rightsquigarrow C$. Since $\models C \equiv C \sqcup \bot$, by (LLE) we get $C \sqcup \bot \rightsquigarrow C$, and then from the definition of $\leq$ follows $C \leq \bot$. ■

**Lemma 17 (∗)** $\prec$ *is a strict partial order on* $\mathcal{D}^{\mathcal{P}}$*, i.e.,* $\prec$ *is irreflexive and transitive.*

**Proof:**
First we show irreflexivity. From the construction of $\prec$, it clearly follows that for every $\langle \mathcal{I}, x, \bot\rangle \in \mathcal{D}^{\mathcal{P}}$, $\langle \mathcal{I}, x, \bot\rangle \not\prec \langle \mathcal{I}, x, \bot\rangle$. Assume that $\langle \mathcal{I}, x, C\rangle \prec \langle \mathcal{I}, x, C\rangle$ for some $C \neq \bot$. Then $C \leq C$ and $x \notin C^{\mathcal{I}}$, i.e., $C \sqcup C \rightsquigarrow C$, and then $C \rightsquigarrow C$, by (LLE). This and $x \notin C^{\mathcal{I}}$ contradicts the fact that $(\mathcal{I}, x)$ is normal for $C$. Hence $\langle \mathcal{I}, x, C\rangle \not\prec \langle \mathcal{I}, x, C\rangle$ for every $\langle \mathcal{I}, x, C\rangle \in \mathcal{D}^{\mathcal{I}}$.

We now show transitivity. Suppose $\langle \mathcal{I}, x, C\rangle \prec \langle \mathcal{I}', x', D\rangle$ and $\langle \mathcal{I}', x', D\rangle \prec \langle \mathcal{I}'', x'', E\rangle$. From the definition of $\prec$ we know that $C, D \neq \bot$, since all non-normal objects are at the highest level in the ordering and are all incomparable. We then have $C \leq D$ and $D \leq E$. (If $E = \bot$, we also have $D \leq E$ by Lemma 16.) From transitivity of $\leq$ (Lemma 10), we conclude $C \leq E$. Since $\langle \mathcal{I}, x, C\rangle \in \mathcal{D}^{\mathcal{P}}$ and $\langle \mathcal{I}, x, C\rangle \prec \langle \mathcal{I}', x', D\rangle$, we conclude that $(\mathcal{I}, x)$ is normal for $C$ and $x \notin D^{\mathcal{P}}$. This and Lemma 14 imply that $x \notin E^{\mathcal{P}}$. ■

**Lemma 18 (∗)** *Given* $\langle \mathcal{I}, x, D\rangle \in \mathcal{D}^{\mathcal{P}}$*,* $\langle \mathcal{I}, x, D\rangle \in \min_{\prec} C^{\mathcal{P}}$ *if and only if* $x \in C^{\mathcal{I}}$ *and* $D \leq C$.

**Proof:**
For the 'if'-part, suppose that $x \in C^{\mathcal{I}}$ and $D \leq C$. Then it clearly follows that $\langle \mathcal{I}, x, D\rangle \in C^{\mathcal{P}}$ (Lemma 12). Now suppose that $\langle \mathcal{I}, x, D\rangle$ is not $\prec$-minimal in $C^{\mathcal{P}}$, i.e., there is $\langle \mathcal{I}', x', E\rangle$ for some $\mathcal{I}'$ such that $x' \in \mathcal{D}^{\mathcal{I}'}$ and some $E \in \mathcal{L}$ such that $\langle \mathcal{I}', x', E\rangle \prec \langle \mathcal{I}, x, D\rangle$ and $x' \in C^{\mathcal{I}'}$. From this and the definition of $\prec$, it follows that $E \leq D$ and $x' \notin D^{\mathcal{I}'}$. Hence $E \leq D \leq C$ and $(\mathcal{I}', x')$ is normal for $E$, and since $x' \in C^{\mathcal{I}'}$, by Lemma 14 we get that $(\mathcal{I}', x')$ is normal for $D$, from which we conclude $x' \in D^{\mathcal{I}'}$, a contradiction.

For the 'only-if'-part, suppose that $\langle \mathcal{I}, x, D\rangle$ is $\prec$-minimal in $C^{\mathcal{P}}$. Then clearly $x \in C^{\mathcal{I}}$. Now assume that there is some $(\mathcal{I}', x')$ which is normal for $C \sqcup D$ and $x' \notin D^{\mathcal{I}'}$. Since $C \sqcup D \leq D$, we must have $\langle \mathcal{I}', x', C \sqcup D\rangle \prec \langle \mathcal{I}, x, D\rangle$. Since $(\mathcal{I}', x')$ is normal for $C \sqcup D$ and $x' \notin D^{\mathcal{I}'}$, it follows that $x' \in C^{\mathcal{I}'}$. This contradicts the minimality of $\langle \mathcal{I}, x, D\rangle$ in $C^{\mathcal{P}}$. Hence all normal $(\mathcal{I}', x')$ for $C \sqcup D$ must satisfy $D$. From this and Lemma 8 follows $C \sqcup D \rightsquigarrow D$, i.e., $D \leq C$.
■

**Lemma 19** **(∗)** *There is no $C \in \mathcal{L}$ such that $C^{\mathcal{P}} \neq \emptyset$ and $\bot \leq C$.*

**Proof:**
Let $C \in \mathcal{L}$ be such that $C^{\mathcal{P}} \neq \emptyset$. Assume that $\bot \leq C$. Then $\bot \sqcup C \rightsquigarrow \bot$, i.e., $C \rightsquigarrow \bot$. Then $C \in \mathcal{C}$, and then $C^{\mathcal{P}} = \emptyset$ by the construction of $\mathcal{P}$. ■

**Corollary 3** **(∗)** *It follows from the two last lemmas that there is no $C \in \mathcal{L}$ for which any $\langle \mathcal{I}, x, \bot\rangle \in \mathcal{D}^{\mathcal{P}}$ is minimal.*

**Lemma 20** **(∗)** *For any $C \in \mathcal{L}$, $C^{\mathcal{P}}$ is smooth.*

**Proof:**
Suppose that $\langle \mathcal{I}, x, D\rangle \in C^{\mathcal{P}}$, i.e., $x \in C^{\mathcal{I}}$. If $D \leq C$, then by Lemma 18 $\langle \mathcal{I}, x, D\rangle$ is $\prec$-minimal in $C^{\mathcal{P}}$. On the other hand, i.e., if $D \not\leq C$, $C \sqcup D \not\rightsquigarrow D$, then by Lemma 8 there is a normal $(\mathcal{I}', x')$ for $C \sqcup D$ such that $x \notin D^{\mathcal{I}'}$. But $C \sqcup D \rightsquigarrow C \sqcup D$, and then $(C \sqcup D) \sqcup D \rightsquigarrow C \sqcup D$, and then $C \sqcup D \leq D$. Hence $\langle \mathcal{I}', x', C \sqcup D\rangle \prec \langle \mathcal{I}, x, D\rangle$. But $x' \in (C \sqcup D)^{\mathcal{I}'}$ and $x' \notin D^{\mathcal{I}'}$, therefore $x' \in C^{\mathcal{I}'}$. Since $C \sqcup D \leq C$, from Lemma 18 we conclude that $\langle \mathcal{I}', x', C \sqcup D\rangle$ is $\prec$-minimal in $C^{\mathcal{P}}$. ■

Next we show in Lemma 21 that the abstract relation $\rightsquigarrow$ we started off with coincides with the relation $\rightsquigarrow_{\mathcal{P}}$ obtained from our constructed preferential interpretation $\mathcal{P}$.

**Lemma 21** **(∗)** *$C \rightsquigarrow D$ if and only if $C \rightsquigarrow_{\mathcal{P}} D$.*

**Proof:**
For the 'only if'-part, we show that $\min_{\prec} C^{\mathcal{P}} \subseteq D^{\mathcal{P}}$. Let $\langle \mathcal{I}, x, E\rangle$ be $\prec$-minimal in $C^{\mathcal{P}}$. Then $(\mathcal{I}, x)$ is normal for $E$ and $x \in C^{\mathcal{P}}$, and from Lemma 18 we also have $E \leq C$. From these results and Lemma 12 it follows that $(\mathcal{I}, x)$ is normal for $C$. Since $C \rightsquigarrow D$, we have $x \in D^{\mathcal{I}}$, and therefore $\langle \mathcal{I}, x, E\rangle \in \mathcal{D}^{\mathcal{P}}$.

Let $C \rightsquigarrow_{\mathcal{P}} D$. From the definition of $\prec$, it follows that for every $(\mathcal{I}, x)$ normal for $C$, $\langle \mathcal{I}, x, C\rangle \in \min_{\prec} C^{\mathcal{P}}$. Since $C \rightsquigarrow_{\mathcal{P}} D$, then $y \in D^{\mathcal{I}'}$ for every $(\mathcal{I}', y)$ that is normal for $C$. This and Lemma 8 give us $C \rightsquigarrow D$. ■

**Proof of Theorem 1**:
Soundness, the 'if'-part, is given in Section A.1. For the 'only if'-part, let $\rightsquigarrow$ be a preferential subsumption and let $\mathcal{P}$ be defined as above. Lemmas 15, 17 and 20 show that $\mathcal{P}$ is a preferential DL interpretation. Lemma 21 shows that $\mathcal{P}$ defines a conditional that is exactly $\rightsquigarrow$. ■

### A.2 Proof Sketch of Theorem 2

Satisfaction of the basic KLM properties for preferential subsumption follows from the proof in Section A.1, given the fact that ranked interpretations are a special case of preferential interpretations. Below we show that rational monotonicity is satisfied.

Assume that $C \leadsto_{\mathcal{R}} E$ but $C \not\leadsto_{\mathcal{R}} \neg D$. From the latter it follows that there is $x \in \min_{\prec} C^{\mathcal{R}}$ such that $x \in D^{\mathcal{R}}$, i.e., $x \in (C \sqcap D)^{\mathcal{R}}$. Let now $x' \in \min_{\prec}(C \sqcap D)^{\mathcal{R}}$. Since $x \in (C \sqcap D)^{\mathcal{R}}$, $x \not\prec x'$ and then $rk(x') < rk(x)$. This means that $x' \in \min_{\prec} C^{\mathcal{R}}$, for if there is $x''$ such that $x'' \prec x'$, then $rk(x'') \leq rk(x')$ and therefore $rk(x'') < rk(x)$ and $x'' \prec x$, which is impossible since $x$ is minimal in $C^{\mathcal{R}}$. From $x' \in \min_{\prec} C^{\mathcal{R}}$ and $\mathcal{R} \Vdash C \precsim E$ follows $x' \in E^{\mathcal{R}}$. Hence $\mathcal{R} \Vdash C \sqcap D \precsim E$ and therefore $C \sqcap D \leadsto_{\mathcal{R}} E$.

The proof of the completeness part relies on the results for the preferential case (Section A.1), with the main difference being the definition of the preference relation, which is shown to be a smooth modular order. This ensures that the canonical model constructed in the proof is a ranked interpretation. Below we provide a sketch of the proof.

Let $\leadsto \subseteq \mathcal{L} \times \mathcal{L}$ satisfy all the basic properties of preferential subsumption together with rational monotonicity.

**Lemma 22** $(*)$ *If $\leadsto$ is rational, then the property below holds:*

$$\frac{C \sqcup E \leadsto \neg C, \;\; D \sqcup E \not\leadsto \neg D}{C \sqcup D \leadsto \neg C}$$

**Definition 27** Let $C \in \mathcal{L}$. We say that $C$ is consistent w.r.t. $\leadsto$ iff $C \not\leadsto \bot$. Given $\mathcal{R} = \langle \mathcal{D}^{\mathcal{R}}, \cdot^{\mathcal{R}}, \prec \rangle$, we say that $C$ is consistent w.r.t. $\leadsto \mathcal{R}$ iff $C \not\leadsto_{\mathcal{R}} \bot$, i.e., iff there is $x \in \mathcal{D}^{\mathcal{R}}$ s.t. $x \in C^{\mathcal{R}}$.

Let $\mathcal{C} = \{C \in \mathcal{L} \mid C \text{ is consistent w.r.t. } \leadsto\}$.

**Lemma 23** $(*)$ *Let $C \in \mathcal{L}$ and let $\leadsto$ be a rational relation. Then $C$ is consistent w.r.t. $\leadsto$ iff there is $(\mathcal{I}, x) \in \mathscr{U}$ s.t. $(\mathcal{I}, x)$ is normal for $C$.*

**Definition 28** Given $C, D \in \mathcal{C}$, $C$ is not more exceptional than $D$, written $C \mathscr{R} D$, iff $C \sqcup D \not\leadsto \neg C$. We say that $C$ is as exceptional as $D$, written $C \sim D$, iff $C \mathscr{R} D$ and $D \mathscr{R} C$.

**Lemma 24** $(*)$ *$\mathscr{R}$ is reflexive and transitive.*

That $\sim$ is an equivalence relation follows from the fact that $\mathscr{R}$ is reflexive and transitive (Lemma 24). With $[C]$ we denote the equivalence class of $C$. The set of equivalence classes of concepts of $\mathcal{C}$ under $\sim$ is denoted by $[\mathcal{C}]$. We write $[C] \leq [D]$ iff $C \mathscr{R} D$, and $[C] < [D]$ iff $[C] \leq [D]$ and $C \not\sim D$.

**Lemma 25** $(*)$ *The relation $<$ is a strict order on $[\mathcal{C}]$.*

**Lemma 26** $(*)$ *Let $C, D \in \mathcal{L}$ be consistent w.r.t. $\leadsto$. If $[C] < [D]$, then $C \leadsto \neg D$.*

**Lemma 27** **($*$)** *Let $C, D \in \mathcal{L}$ be consistent w.r.t. $\rightsquigarrow$. If there is $(\mathcal{I}, x) \in \mathscr{U}$ s.t. $(\mathcal{I}, x)$ is normal for $C$ and $x \in D^{\mathcal{I}}$, then $[D] \leq [C]$.*

Armed with these results, we can then construct an interpretation $\mathcal{R}$ analogous to the preferential interpretation $\mathcal{P}$ in Section A.1, with the preference relation defined as follows:

- For every $\langle \mathcal{I}, x, C \rangle \in \mathcal{D}^{\mathcal{R}}$ such that $C \neq \bot$, $\langle \mathcal{I}, x, C \rangle \prec \langle \mathcal{J}, y, \bot \rangle$ for every $\langle \mathcal{J}, y, \bot \rangle \in \mathcal{D}^{\mathcal{R}}$;
- For every $\langle \mathcal{I}, x, C \rangle, \langle \mathcal{J}, y, D \rangle \in \mathcal{D}^{\mathcal{P}}$ such that $C, D \neq \bot$, $\langle \mathcal{I}, x, C \rangle \prec \langle \mathcal{J}, y, D \rangle$ if and only if $[C] < [D]$.

**Lemma 28** **($*$)** *$\prec$ is a modular partial order.*

**Lemma 29** **($*$)** *For every $C \in \mathcal{L}$, if $C$ is consistent, then $C^{\mathcal{R}}$ is smooth.*

From this point on, a result analogous to Lemma 21 above can be shown to hold for the defeasible conditional $\rightsquigarrow_{\mathcal{R}}$ induced by $\mathcal{R}$. From that the result follows. ∎

## Appendix B. Proof outlines for Section 4

The proofs in this section are mostly traightforward generalisations, or easy corollaries, of results proved by Lehmann and Magidor (1992) for the propositional case. In order to get to these, we first provide some supporting results.

**Lemma 30** **($*$)** *Consider a preferential interpretation $\mathcal{P}$, pick any $E \in \mathcal{L}$ and let $x \in \min_{\prec^{\mathcal{P}}}(E^{\mathcal{P}})$. Let $\prec^{\mathcal{P}}_{(E,x)}$ be the strict partial order obtained by making $x$ the $\prec^{\mathcal{P}}$-minimum of $E^{\mathcal{P}}$. That is, $u \prec^{\mathcal{P}}_{(E,x)} v$ iff $u \prec^{\mathcal{P}} v$ or $u \preceq^{\mathcal{P}} x$ and there is a $y \in E^{\mathcal{P}}$ s.t. $y \preceq^{\mathcal{P}} v$. Let $\mathcal{P}^{(E,x)}$ be the structure $\langle \Delta^{\mathcal{P}}, \cdot^{\mathcal{P}}, \prec^{\mathcal{P}}_{(E,x)} \rangle$. Then $\mathcal{P}^{(E,x)}$ is a preferential interpretation in which $x$ is a $\prec^{\mathcal{P}^{(E,x)}}$-minimum of $E^{\mathcal{P}^{(E,x)}}$. If $\mathcal{P} \Vdash C \sqsubsim D$ then $\mathcal{P}^{(E,x)} \Vdash C \sqsubsim D$. For $C \in \mathcal{L}$, $\mathcal{P} \Vdash C \sqsubsim \bot$ iff $\mathcal{P}^{(E,x)} \Vdash C \sqsubsim \bot$*

The proof is similar to that of Lemma 2.17 on page 11 by Lehmann and Magidor (1992).

**Lemma 31** **($*$)** *If $\mathcal{K} \not\models_P C \sqsubsim D$ then $\mathcal{K} \models_P E \sqsubsim \bot$ iff $\mathcal{K} \cup \{C \sqsubsim \neg D\} \models_P E \sqsubsim \bot$.*

The proof is similar to that of Theorem 2.18 on page 11 by Lehmann and Magidor (1992)—it uses Lemma 30.

**Proposition 1** *If a knowledge base $\mathcal{K}$ is preferentially satisfiable, then $\rightsquigarrow^{\mathcal{K}}_{P}$ is preferential.*

**Proof:**
To prove this we need to show that $\rightsquigarrow^{\mathcal{K}}_{P}$ satisfies (Cons), (LLE), (And), (Or), (RW), and (CM). (Cons) follows immediately from the fact that $\mathcal{K}$ is preferentially satisfiable. The remainder is tedious, but easy to prove, and is similar to that of Theorem 2.9 on page 9 by

Lehmann and Magidor (1992). ■

**Proposition 2** *Let $\mathrel{\leadsto}_P^{\mathcal{K}}$ be the (preferential) conditional induced by a KB $\mathcal{K}$ under $\models_P$. Then $\mathrel{\leadsto}_P^{\mathcal{K}}$ coincides with the preferential closure of $\mathcal{K}$.*

**Proof:**
Similar to that of Theorem 2.9 on page 9 by Lehmann and Magidor (1992), which refers to a prior result by Kraus et al. (1990). ■

**Proposition 3** *A knowledge base $\mathcal{K}$ is preferentially satisfiable if and only if it is rank satisfiable.*

**Proof:**
The *if* part follows immediately. For the converse, suppose there is preferential interpretation $\mathcal{P}$ s.t. $\mathcal{P} \Vdash \mathcal{K}$. By Theorem 1, $\leadsto_P$ is preferential. We remark that there exists a rational extension of $\leadsto_P$, say $\leadsto$, s.t. for every $C \in \mathcal{L}$, $C \leadsto_p \bot$ if and only if $C \leadsto \bot$. Proving this is similar to the proof of Lemma 4.1 on page 25 by Lehmann and Magidor (1992). The proof makes use of Lemma 31. It is worth noting that Lemma 4.1 does not show directly that $\leadsto$ satisfies (Cons), and is therefore rational. To show this, observe that $\top \not\leadsto_P \bot$ and therefore that $\top \not\leadsto \bot$.

By Theorem 2 it follows that there is a ranked interpretation $\mathcal{R}$ such that, if $\mathcal{P} \Vdash \alpha$ then $\mathcal{R} \Vdash \alpha$. So $\mathcal{R} \Vdash \mathcal{K}$, which means $\mathcal{K}$ is rank satisfiable. ■

**Theorem 3** *A subsumption statement $\alpha$ is preferrentially entailed by a knowledge base $\mathcal{K}$ if and only if it is rank entailed by $\mathcal{K}$. That is,*

$$\mathcal{K} \models_P \alpha \text{ iff } \mathcal{K} \models_R \alpha.$$

The proof is similar to that of Theorem 4.2 on page 25 by Lehmann and Magidor (1992).

## Appendix C. Proofs for Section 5

**Proposition 4** *For a KB $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$ and a defeasible axiom $C \mathrel{\sqsubseteq\!\!\!\sim} D$, $\mathcal{K} \models_R C \mathrel{\sqsubseteq\!\!\!\sim} D$ if and only if for every $R \in R_\Delta^{\mathcal{K}}$, $R \Vdash C \mathrel{\sqsubseteq\!\!\!\sim} D$.*

**Proof:**
Let $\Delta$ be a countably infinite domain. If $\mathcal{K} \models_R C \mathrel{\sqsubseteq\!\!\!\sim} D$, then obviously $\forall R \in R_\Delta^{\mathcal{K}}$, $R \Vdash C \mathrel{\sqsubseteq\!\!\!\sim} D$. Assume $\forall R \in R_\Delta^{\mathcal{K}}$, $R \Vdash C \mathrel{\sqsubseteq\!\!\!\sim} D$. We have to prove that it is not possible that there is some $R$ in $R^{\mathcal{K}}$ s.t. $R \not\Vdash C \mathrel{\sqsubseteq\!\!\!\sim} D$. Let $R$ be a model of $\mathcal{K}$ and a counter-model of $C \mathrel{\sqsubseteq\!\!\!\sim} D$. We can easily prove that for a defeasible language built over the description logic $\mathcal{ALC}$ (or any of $\mathcal{ALC}$-sublanguages) and semantically characterised by the class of the ranked models, the FMP (Finite Model Property) holds (see Appendix G). Then, if we have a ranked interpretation $R$ that is a model of $\mathcal{K}$ and a counter-model of $C \mathrel{\sqsubseteq\!\!\!\sim} D$ (with a domain that could be countable or uncountable), there must be a model $R_{fin}$ with a finite domain that is too a model of $\mathcal{K}$ and a counter-model of $C \mathrel{\sqsubseteq\!\!\!\sim} D$. Given $R_{fin}$, then we can extend it to a model of $\mathcal{K}$ that

is a counter-model of $C \precsim D$ with a countably infinite domain. Hence, if there is a counter model of $\mathcal{K} \models_R C \precsim D$, there must be also a counter model with a countable domain, and we can consider only the models with a countably infinite domain in the definition of $\models_R$.

Now, let $R' = \langle \Delta', \cdot^{R'}, \prec^{R'} \rangle$ be a model of $\mathcal{K}$ and a counter-model of $C \precsim D$ with $\Delta'$ countably infinite. It is easy to build an isomorphic interpretation $R = \langle \Delta, \cdot^{R}, \prec^{R} \rangle$: once we have defined a bijection $b : \Delta' \times \Delta$ (that must exist, being both $\Delta'$ and $\Delta$ countably infinite sets), we can define $\cdot^R$ and $\prec^R$ in the following way:

- For every $r \in \mathsf{N}_{\mathscr{R}}$ and every $x, y \in \Delta'$, $\langle b(x), b(y) \rangle \in r^R$ iff $\langle x, y \rangle \in r^{R'}$;
- For every $A \in \mathsf{N}_{\mathscr{C}}$ and every $x \in \Delta'$, $b(x) \in A^R$ iff $x \in A^{R'}$;
- For every $x, y \in \Delta'$, $b(x) \prec^R b(y)$ iff $x \prec^{R'} y$.

It is easy to prove by induction on the construction of the concepts that for every concept $C$, $x \in C^{R'}$ iff $b(x) \in C^R$; moreover, $x \in \min_{\prec^{R'}}(C^{R'})$ iff $b(x) \in \min_{\prec^R}(C^R)$. Hence, if there is a countermodel to $\mathcal{K} \models_R C \precsim D$ there must be also a counter model with $\Delta$ as domain. Hence, we can use just the set of interpretations in $R^{\mathcal{K}}_{\Delta}$ to decide the consequences of $\mathcal{K}$ w.r.t. rank entailment. ■

**Theorem 4** *Let $\mathcal{K}$ be a knowledge base having a ranked model. Then $R^{\mathcal{K}}_{\cup}$ is a model of $\mathcal{K}$ and for any pair of concepts $C, D$, $R^{\mathcal{K}}_{\cup} \Vdash C \precsim D$ if and only if $r_{\mathcal{K}}(C \sqcap D) < r_{\mathcal{K}}(C \sqcap \neg D)$ or $r_{\mathcal{K}}(C) = \infty$.*

**Proof:**
First of all we have to prove that the exceptionality function of Definition 12 is correctly represented in this model, that is, $R^{\mathcal{K}}_{\cup} \Vdash \top \precsim \neg C$ iff $\mathcal{K} \models_R \top \precsim \neg C$. By Proposition 4, a concept $C$ is exceptional w.r.t. $\mathcal{K}$ iff

$$\forall R \in R^{\mathcal{K}}_{\Delta}, R \Vdash \top \precsim \neg C$$

that immediately corresponds to saying

$$R^{\mathcal{K}}_{\cup} \Vdash \top \precsim \neg C.$$

Now we have to prove that this correspondence is preserved for all the steps of the exceptionality function, that is, we have to prove that the model $R^{\mathcal{K}}_{\cup}$ is a correct semantical representation of the ranking procedure. That corresponds to saying that for every concept $C$ and every $i$, $0 < i \leq n$,

$$r_{\mathcal{K}}(C) = i \text{ iff } h_{R^{\mathcal{K}}_{\cup}}(C) = i$$

We can prove it by induction on the ranking value $i$ ($i > 0$).

If $h_{R^{\mathcal{K}}_{\cup}}(C) = i$ it is immediate that $r_{\mathcal{K}}(C) \leq i$. We have to prove that if $r_{\mathcal{K}}(C) = i$ then $h_{R^{\mathcal{K}}_{\cup}}(C) = i$. We can prove that by defining a model $R$ in $R^{\mathcal{K}}$ s.t. $h_R(C) = i$.

So, given a KB $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$, let $r_{\mathcal{K}}(C) = i$, and let $\mathcal{D}^r_{\geq i}$ be the subset of $\mathcal{D}$ containing the defeasible axioms with a ranking value of at least $i$. Let $M$ be a ranked model of $\langle \mathcal{T}, \mathcal{D}^r_{\geq i} \rangle$ s.t. $h_M(C) = 0$; such a model must exist, since $r_{\mathcal{K}}(C) = i$, that is, $C$ is not exceptional in $\langle \mathcal{T}, \mathcal{D}^r_{\geq i} \rangle$. We assume that $M$ has a finite domain (we can, due to the FMP).

Now, let $N$ be a model of $\langle \mathcal{T}, \mathcal{D}\rangle$ in $R^{\mathcal{K}}_{\Delta}$ s.t. for all the axioms $C \precsim D \in (\mathcal{D} \setminus \mathcal{D}^{r}_{\geq i})$ there is an individual satisfying $C \sqcap D$. The induction hypothesis guarantees that such a model exists.

We define a new interpretation $N' = \langle \Delta^{N'}, \cdot^{N'}, \prec^{N'}\rangle$ in the following way:

- $\Delta^{N'} = \Delta^{M} \cup \Delta^{N}$;
- for every atomic concept $A$ and every $a \in \Delta^{N'}$, $a \in A^{N'}$ iff one of the two following cases holds: either $a \in \Delta^{N}$ and $a \in A^{N}$, or $a \in \Delta^{M}$ and $a \in A^{M}$;
- for every role $r$ and every $a, b \in \Delta^{N'}$, $(a, b) \in r^{N'}$ iff one of the two following cases holds: either $(a, b) \in \Delta^{N}$ and $(a, b) \in r^{N}$, or $(a, b) \in \Delta^{M}$ and $(a, b) \in r^{M}$;
- For every $a \in \Delta^{N'}$, $h_{N'}(a) = j$ iff one of the two following cases holds: either $a \in \Delta^{N}$ and $h_{N}(a) = j$, or $a \in \Delta^{M}$ and $h_{M}(a) = j - i$.

and so on, until we finish the layers of both $N$ and $M$. $N'$ is a model of $\langle \mathcal{T}, \mathcal{D}\rangle$ (easy to prove by induction on the construction of the concepts) s.t. $h_{N'}(C) = i$. Since $N'$ is obtained from the composition of a model with $\Delta$ as domain and a model with a finite domain, that is, $N'$ has a countably infinite domain, there is a model $N'_{\Delta}$ of $\langle \mathcal{T}, \mathcal{D}\rangle$ that is isomorphic to $N'$ and has $\Delta$ as domain. So $N'_{\Delta}$ takes part in the construction of $R^{\mathcal{K}}_{\cup}$, and $h_{R^{\mathcal{K}}_{\cup}}(C) = i$.

This proves that for every $C$, $r_{\mathcal{K}}(C) = i$ iff $h_{R^{\mathcal{K}}_{\cup}}(C) = i$.

Since $R^{\mathcal{K}}_{\cup} \Vdash C \precsim D$ iff $h_{R^{\mathcal{K}}_{\cup}}(C \sqcap D) < h_{R^{\mathcal{K}}_{\cup}}(C \sqcap \neg D)$ (or $h_{R^{\mathcal{K}}_{\cup}}(C) = \infty$), and $h_{R^{\mathcal{K}}_{\cup}}(C \sqcap D) < h_{R^{\mathcal{K}}_{\cup}}(C \sqcap \neg D)$ (or $h_{R^{\mathcal{K}}_{\cup}}(C) = \infty$) iff $r_{\mathcal{K}}(C \sqcap D) < r_{\mathcal{K}}(C \sqcap \neg D)$ (or $r(C) = \infty$), we have proven the proposition.

■

## Appendix D. Proofs for Section 6

**Proposition 5** *Given a knowledge base $\mathcal{K} = \langle \mathcal{T}, \mathcal{D}\rangle$, $\mathcal{T} \cup \mathcal{D} \subseteq C_r(\langle \mathcal{T}, \mathcal{D}\rangle)$. Moreover, $C_r(\mathcal{K})$ defines a defeasible conditional $\rightsquigarrow^{\mathcal{K}}_{r}$ that is rational, where $\rightsquigarrow^{\mathcal{K}}_{r} := \{(C, D) \mid \mathcal{K} \vdash_r C \precsim D\}$.*

**Proof:**
Assume that $C \sqsubseteq D \in \mathcal{T}$. $\langle \mathcal{T}, \mathcal{D}\rangle \vdash_r C \sqsubseteq D$ iff $\mathcal{T}^* \models C \sqsubseteq D$; since $\mathcal{T} \subseteq \mathcal{T}^*$, $\mathcal{T} \subseteq C_r(\langle \mathcal{T}, \mathcal{D}\rangle)$.

Assume that $C \precsim D \in \mathcal{D}$. Either $C \precsim D$ ends up in $\mathcal{D}^*_{\infty}$, or there will be an $i$ ($0 \leq i \leq n$) s.t. $r(C) = r(C \precsim D) = i$. In the former case, $C \sqsubseteq D$ is in $\mathcal{T}^*$, and so $\mathcal{T}^* \models C \sqsubseteq D$, i.e., $\langle \mathcal{T}, \mathcal{D}\rangle \vdash_r C \precsim D$. In the latter case, $\models \overline{\mathcal{E}_i} \sqsubseteq \neg C \sqcup D$, and so $\mathcal{T}^* \models \overline{\mathcal{E}_i} \sqcap C \sqsubseteq D$, i.e., $C \precsim D \in C_r(\langle \mathcal{T}, \mathcal{D}\rangle)$. Hence $\mathcal{T} \cup \mathcal{D} \subseteq C_r(\langle \mathcal{T}, \mathcal{D}\rangle)$.

Given Definition 14, to prove that $C_r(\langle \mathcal{T}, \mathcal{D}\rangle)$ satisfies the rational properties w.r.t. $\precsim$ is quite straightforward.

- (Ref). Since $\models C \sqsubseteq C$ is valid in $\mathcal{ALC}$ for any $C$, we have that $\mathcal{T}^* \models \overline{\mathcal{E}_i} \sqcap C \sqsubseteq C$ for any $\mathcal{T}^*$ and $\overline{\mathcal{E}_i}$.
- (LLE). $C \precsim E \in C_r(\langle \mathcal{T}, \mathcal{D}\rangle)$ implies that $\mathcal{T}^* \models \overline{\mathcal{E}_i} \sqcap C \sqsubseteq E$ for some $i$ (or $\mathcal{T}^* \models C \sqsubseteq E$, if $r(C) = \infty$). Since $\models C \equiv D$, $\overline{\mathcal{E}_i}$ is the lowest $i$ s.t. $\mathcal{T}^* \not\models \bigsqcap \overline{\mathcal{E}_i} \sqsubseteq \neg D$, and $\mathcal{T}^* \models \overline{\mathcal{E}_i} \sqcap D \sqsubseteq E$ too.

- (And). $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_i} \sqcap C \sqsubseteq D$ and $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_i} \sqcap C \sqsubseteq E$ (possibly without the $\bigsqcap \overline{\mathcal{E}_i}$, if $C$ has an infinite rank), hence $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_i} \sqcap C \sqsubseteq D \sqcap E$, that is, $C \precsim D \sqcap E \in C_r(\langle \mathcal{T}, \mathcal{D} \rangle)$.

- (Or). $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_i} \sqcap C \sqsubseteq E$ for some $i$ and $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_j} \sqcap D \sqsubseteq E$ for some $j$. Let's assume that $i \leq j$ and $i < \infty$, that is, $\models \bigsqcap \overline{\mathcal{E}_i} \sqsubseteq \bigsqcap \overline{\mathcal{E}_j}$. Then, since $\mathcal{T}^* \not\models \bigsqcap \overline{\mathcal{E}_i} \sqsubseteq \neg C$, we have that $\mathcal{T}^* \not\models \bigsqcap \overline{\mathcal{E}_i} \sqsubseteq \neg (C \sqcup D)$. Moreover $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_j} \sqcap D \sqsubseteq E$ and $\models \bigsqcap \overline{\mathcal{E}_i} \sqsubseteq \bigsqcap \overline{\mathcal{E}_j}$ imply that $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_i} \sqcap D \sqsubseteq E$. So, $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_i} \sqcap (C \sqcup D) \sqsubseteq E$. The proof is analogous for $j \leq i$ with $j < \infty$, or if $i$ and $j$ correspond to $\infty$.

- (RW). $C \precsim D \in C_r(\langle \mathcal{T}, \mathcal{D} \rangle)$ if $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_i} \sqcap C \sqsubseteq D$ for some $\bigsqcap \overline{\mathcal{E}_i}$ (or $\mathcal{T}^* \models C \sqsubseteq D$, if $r(C) = \infty$). Since $\models D \sqsubseteq E$, $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_i} \sqcap C \sqsubseteq E$.

- (CM). If $r(C) = i < \infty$, $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_i} \sqcap C \sqsubseteq D$ and $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_i} \sqcap C \sqsubseteq E$ for some $\bigsqcap \overline{\mathcal{E}_i}$. Since $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_i} \sqcap C \sqsubseteq D$ and $\mathcal{T}^* \not\models \bigsqcap \overline{\mathcal{E}_i} \sqsubseteq \neg C$, $\mathcal{T}^* \not\models \bigsqcap \overline{\mathcal{E}_i} \sqsubseteq \neg (C \sqcap D)$, otherwise we would have $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_i} \sqcap C \sqsubseteq D \sqcap \neg D$, i.e., $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_i} \sqsubseteq \neg C$. Hence we have $C \sqcap D \precsim E \in C_r(\langle \mathcal{T}, \mathcal{D} \rangle)$ since $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_i} \sqcap C \sqcap D \sqsubseteq E$. If $r(C) = \infty$, we have $\mathcal{T}^* \models C \sqsubseteq \bot$, and the proof is trivial.

- (RM). If $r(C) = i < \infty$, $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_i} \sqcap C \sqsubseteq E$ and $\mathcal{T}^* \not\models \bigsqcap \overline{\mathcal{E}_i} \sqcap C \sqsubseteq \neg D$ for some $\bigsqcap \overline{\mathcal{E}_i}$. Since $\mathcal{T}^* \not\models \bigsqcap \overline{\mathcal{E}_i} \sqcap C \sqsubseteq \neg D$ and $\mathcal{T}^* \not\models \bigsqcap \overline{\mathcal{E}_i} \sqsubseteq \neg C$, $\mathcal{T}^* \not\models \bigsqcap \overline{\mathcal{E}_i} \sqsubseteq \neg (C \sqcap D)$, otherwise we would have $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_i} \sqcap C \sqsubseteq \neg D$. Hence we have $C \sqcap D \precsim E \in C_r(\langle \mathcal{T}, \mathcal{D} \rangle)$ since $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_i} \sqcap C \sqcap D \sqsubseteq E$. If $r(C) = \infty$, we have $\mathcal{T}^* \models C \sqsubseteq \bot$, and the proof is trivial.

■

**Lemma 2** *Assume that $C \precsim D \in \mathcal{D}$. $\langle \mathcal{T}, \mathcal{D} \rangle \models_R C \precsim \bot$ if and only if $r(C) = \infty$, i.e., if and only if $\mathcal{T}^* \models C \sqsubseteq \bot$.*

**Proof:**

From left to right: $\langle \mathcal{T}, \mathcal{D} \rangle \models_R C \precsim \bot$ implies that every rational subsumption relation containing $\langle \mathcal{T}, \mathcal{D} \rangle$ must satisfy also $C \precsim \bot$, hence we have that $\langle \mathcal{T}, \mathcal{D} \rangle \vdash_r C \precsim \bot$, since $C_r(\langle \mathcal{T}, \mathcal{D} \rangle)$ defines a rational subsumption relation satisfying $\langle \mathcal{T}, \mathcal{D} \rangle$ (Proposition 5). From Definition 14 we know that $\langle \mathcal{T}, \mathcal{D} \rangle \vdash_r C \precsim \bot$ is possible only if $C$ is always negated in the ranking procedure, i.e., $\mathcal{T}^* \models C \sqsubseteq \bot$.

From right to left: We define from $\langle \mathcal{T}, \mathcal{D} \rangle$ a new knowledge base $\langle \mathcal{T}^*, \mathcal{D}^* \rangle$, with $\mathcal{T}^*$ obtained from $\mathcal{T}$ adding all the sets $\{C \sqsubseteq D \mid C \precsim D \in \mathcal{D}^*_\infty\}$ that we obtain at each reiteration of Algorithm ComputeRanking. Let's indicate with $\mathcal{D}^1_\sqsubseteq, \ldots, \mathcal{D}^n_\sqsubseteq$ such sets. Assume that $\mathcal{T}^* \models C \sqsubseteq \bot$, but $\langle \mathcal{T}, \mathcal{D} \rangle \not\models_R C \precsim \bot$, i.e., there is a ranked model of $\langle \mathcal{T}, \mathcal{D} \rangle$ s.t. $C$ is non-empty. Consider such a model $\mathcal{M}$, with an object $a$ falling under $C^\mathcal{I}$. Since $\mathcal{T}^* \models C \sqsubseteq \bot$, there must be a subsumption axiom $E \sqsubseteq F$ in some $\mathcal{D}^i_\sqsubseteq$ that is not satisfied, that, given the nature of the axioms in every $\mathcal{D}^n_\sqsubseteq$ ($\mathcal{T}^* \models E \sqsubseteq \bot$ for every $E \sqsubseteq F$ contained in some $\mathcal{D}^n_\sqsubseteq$), means that there is a subsumption $E \sqsubseteq \bot$ that is not satisfied in $\mathcal{M}$. Hence there must be an individual $b$ falling under $E$ in $\mathcal{M}$. Hence, assuming $E \sqsubseteq F \in \mathcal{D}^i_\sqsubseteq$, since $\mathcal{T} \cup \mathcal{D}^1_\sqsubseteq \cup \ldots \cup \mathcal{D}^{i-1}_\sqsubseteq \models \bigsqcap\{\neg G \sqcup H \mid G \sqsubseteq H \in \mathcal{D}^i_\sqsubseteq\} \sqsubseteq \neg E$, either $\mathcal{M} \Vdash \mathcal{T} \cup \mathcal{D}^1_\sqsubseteq \cup \ldots \cup \mathcal{D}^{i-1}_\sqsubseteq$ and $\mathcal{M} \Vdash b : G \sqcap \neg H$ for some $G \sqsubseteq H \in \mathcal{D}^i_\sqsubseteq$ (Case 1), or $\mathcal{M} \not\Vdash \mathcal{T} \cup \mathcal{D}^1_\sqsubseteq \cup \ldots \cup \mathcal{D}^{i-1}_\sqsubseteq$ (Case 2).

Case 1. Since $\mathcal{M}$ is a model of $\langle \mathcal{T}, \mathcal{D} \rangle$, hence it is a model also of $G \mathrel{\tilde{\sqsubseteq}} H$, that is in $\mathcal{D}$. Hence there must be an individual $c$ s.t. $c \prec b$ and $c : G \sqcap H$. Again, since $G \sqsubseteq H \in \mathcal{D}^i_{\sqsubseteq}$ (that implies $\mathcal{T} \cup \mathcal{D}^1_{\sqsubseteq} \cup \ldots \cup \mathcal{D}^{i-1}_{\sqsubseteq} \models \bigsqcap\{\neg G \sqcup H \mid G \sqsubseteq H \in \mathcal{D}^i_{\sqsubseteq}\} \sqsubseteq \neg G$) and $\mathcal{M} \Vdash \mathcal{T} \cup \mathcal{D}^1_{\sqsubseteq} \cup \ldots \cup \mathcal{D}^{i-1}_{\sqsubseteq}$, there must be an axiom $I \sqsubseteq L \in \mathcal{D}^i_{\sqsubseteq}$ s.t. $\mathcal{M} \vdash c : I \sqcap \neg L$, and we need an individual $d$ s.t. $d \prec c$ and $\mathcal{M} \vdash d : H \sqcap I$, and so on...

This procedure creates an infinite descending chain of individuals, and, since the number of the antecedents of the axioms in $\mathcal{D}^i_{\infty}$ is finite, it cannot be the case since the model would not satisfy the smoothness condition for the concept $\bigsqcup\{C | C \mathrel{\tilde{\sqsubseteq}} D \in \mathcal{D}^i_{\infty}\}$ (see Definition 1).

Case 2. If $\mathcal{M} \not\Vdash \mathcal{T} \cup \mathcal{D}^1_{\sqsubseteq} \cup \ldots \cup \mathcal{D}^{i-1}_{\sqsubseteq}$, then $\mathcal{M}$ does not satisfy some axiom $E \sqsubseteq F \in \mathcal{D}^j_{\sqsubseteq}$ for some $j < i$, and therefore there must be an object falling under $E$ in $\mathcal{M}$. Again, it is Case 1 or Case 2. However, since at every reiteration of Case 2 we pick a lower value $j$ for $\mathcal{D}^j_{\sqsubseteq}$ and we have a finite sequence of $\mathcal{D}^j_{\sqsubseteq}$, we know that after some steps (in the worst case when we reach $\mathcal{D}^0_{\sqsubseteq}$) we necessarily fall into Case 1, that cannot be the case.

■

**Proposition 6** *Consider a knowledge base $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$ and the knowledge base $\mathcal{K}^* = \langle \mathcal{T}^*, \mathcal{D}^* \rangle$ obtained from $\mathcal{K}$ using Procedure ComputeRanking. $\mathcal{K}$ and $\mathcal{K}^*$ are rank equivalent.*

**Proof:**

We have $\mathcal{T}^* = \mathcal{T} \cup \{C_1 \sqsubseteq D_1, \ldots, C_n \sqsubseteq D_n\}$ and $\mathcal{D}^* = \mathcal{D}/\{C_1 \mathrel{\tilde{\sqsubseteq}} D_1, \ldots, C_n \mathrel{\tilde{\sqsubseteq}} D_n\}$. It is sufficient to prove that $\mathcal{K} \models_R C_i \sqsubseteq \bot$ and $\mathcal{K}^* \models_R C_i \mathrel{\tilde{\sqsubseteq}} D_i$ for every $C_i \mathrel{\tilde{\sqsubseteq}} D_i$ $(1 \leq i \leq n)$.

So, let $C_i \mathrel{\tilde{\sqsubseteq}} D_i \in \mathcal{D} \setminus \mathcal{D}^*$. It means that at some repetition of lines 4-14 of the Procedure ComputeRanking $C_i \mathrel{\tilde{\sqsubseteq}} D_i$ is moved into $of the Procedure Compute Ranking$; that is, in some repetition of the lines 4-14 we have $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{D}^*_{\infty}} \sqsubseteq \neg C_i$, that implies that $\mathcal{T}^* \cup \mathcal{D}^{*\sqsubseteq}_{\infty} \models \top \sqsubseteq \neg C_i$ (where $\mathcal{D}^{*\sqsubseteq}_{\infty} = \{C \sqsubseteq D \mid C \mathrel{\tilde{\sqsubseteq}} D \in \mathcal{D}^*_{\infty}\}$), that in turn implies that, since every $\mathcal{D}^{*\sqsubseteq}_{\infty}$ created at every repetition is contained in the final $\mathcal{T}^*$, using such final $\mathcal{T}^*$ we have that $\mathcal{T}^* \models C \sqsubseteq \bot$; hence, by Lemma 2 we have that $\mathcal{K} \models_R C_i \mathrel{\tilde{\sqsubseteq}} \bot$, i.e., $\mathcal{K} \models_R C_i \sqsubseteq \bot$.

On the other hand, if $C_i \mathrel{\tilde{\sqsubseteq}} D_i \in \mathcal{D} \setminus \mathcal{D}^*$, then $C_i \sqsubseteq D_i \in \mathcal{T}^*$, and hence $\mathcal{K}^* \models_R C_i \mathrel{\tilde{\sqsubseteq}} D_i$ by Supraclassicality (check the proof of Lemma 3). ■

**Lemma 3** *For every concept $C$, $\langle \mathcal{T}, \mathcal{D} \rangle \models_R \top \mathrel{\tilde{\sqsubseteq}} C$ if and only if $\langle \mathcal{T}, \mathcal{D} \rangle \vdash_r \top \mathrel{\tilde{\sqsubseteq}} C$.*

**Proof:**

Remember that $\langle \mathcal{T}, \mathcal{D} \rangle \vdash_r \top \mathrel{\tilde{\sqsubseteq}} C$ iff $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{D}^*} \sqsubseteq C$.

From right to left: First, we need to prove two properties of $\models_R$: supraclassicality (Sup) and one half of the deduction theorem (S):

(Sup):

$$\frac{C \sqsubseteq D}{C \mathrel{\tilde{\sqsubseteq}} D}$$

Assume $C \mathrel{\tilde{\sqsubseteq}} C$ and $C \sqsubseteq D$ and apply (RW).

(S):

$$\frac{C \precsim D}{\top \precsim \neg C \sqcup D}$$

Assume $C \precsim D$ and $\models D \sqsubseteq \neg C \sqcup D$ (that is classically valid); we derive by (RW) $C \precsim \neg C \sqcup D$. Assume $\models \neg C \sqsubseteq \neg C \sqcup D$ (classically valid); we obtain $\neg C \precsim \neg C \sqcup D$ by (Sup). We apply (Or) to $C \precsim \neg C \sqcup D$ and $\neg C \precsim \neg C \sqcup D$, obtaining $\top \precsim \neg C \sqcup D$.

Now we have to prove that if $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{D}^*} \sqsubseteq C$, then $\langle \mathcal{T}, \mathcal{D} \rangle \models_R \top \precsim C$.

From corollary 6 we know that $\mathcal{T}^* \cup \mathcal{D}^*$ is in the ranked consequences of $\langle \mathcal{T}, \mathcal{D} \rangle$. Applying (S) to all the axioms $C \precsim D$ in $\mathcal{D}^*$, and we have $\langle \mathcal{T}, \mathcal{D} \rangle \models_R \top \precsim \neg C \sqcup D$ from each of them. Applying (And) to all these defeasible inclusions, we have $\top \precsim \bigsqcap \overline{\mathcal{D}'}$ and, by (RW'), we obtain $\top \precsim C$.

From left to right: Immediate from Proposition 5.

■

**Corollary 4 (∗)** *For every $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$ and every concept $C$, $r_{\mathcal{K}}(C) = \infty$ iff $r(C) = \infty$.*

**Proof:**
Consider a KB $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$, and transform it into a rank equivalent knowledge base $\mathcal{D}'$ composed of only defeasible axioms (see Lemma 2). Since the model $R_{\cup}^{\mathcal{K}}$ of the RC of $\mathcal{K}$ must be also a model of $\mathcal{D}'$, we can easily derive from Proposition 4 that $\mathcal{K} \models_R^{\cup} C \precsim \bot$ (that is, $r_{\mathcal{K}}(C) = \infty$) iff $\mathcal{K} \models_R C \precsim \bot$. From Lemma 2 we have that $\mathcal{K} \models_R C \precsim \bot$ iff $r(C) = \infty$, hence the result. ■

**Proposition 7** *For every $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$ and every concept $C$, $r_{\mathcal{K}}(C) = r(C)$.*

**Proof:**
From Lemma 2, Corollary 4 and Proposition 6 we can see that, given a knowledge base $\langle \mathcal{T}, \mathcal{D} \rangle$ (possibly, with an empty $\mathcal{T}$), we can define a rank equivalent knowledge base $\langle \mathcal{T}^*, \mathcal{D}^* \rangle$ s.t. all the classical information 'hidden' in $\mathcal{D}$ is moved into $\mathcal{T}^*$. $\langle \mathcal{T}^*, \mathcal{D}^* \rangle$ can be defined identifying the elements of $\mathcal{D}$ that have $\infty$ as ranking value, and Corollary 4 shows that w.r.t. the value $\infty$, $r_{\mathcal{K}}$ and $r$ are equivalent, while Proposition 6 tells us that $\langle \mathcal{T}, \mathcal{D} \rangle$ and $\langle \mathcal{T}^*, \mathcal{D}^* \rangle$ are rank equivalent. Once we have defined $\langle \mathcal{T}^*, \mathcal{D}^* \rangle$, Lemma 3 implies that a concept $C$ is exceptional w.r.t. $\models_R^{\cup}$ ($\mathcal{K} \models_R \top \precsim \neg C$) iff it is exceptional w.r.t. $\vdash_r$. Hence the two ranking functions $r_{\mathcal{K}}$ and $r$ give back exactly the same results. ■

**Theorem 5** *Given a knowledge base $\langle \mathcal{T}, \mathcal{D} \rangle$, for every pair of concepts $C, D$, $\langle \mathcal{T}, \mathcal{D} \rangle \models_R^{\cup} C \precsim D$ iff $\langle \mathcal{T}, \mathcal{D} \rangle \vdash_r C \precsim D$.*

**Proof:**
Since we have proven Proposition 7, in this proof we are generically going to use $r$ to indicate indifferently the equivalent ranking functions $r_{\mathcal{K}}$ and $r$.

From left to right: Assume $\langle \mathcal{T}, \mathcal{D} \rangle \models_R^{\cup} C \precsim D$. That means that either $r(C \sqcap \neg D) > r(C)$ or $r(C) = \infty$. In the first case, it means that there is an $i$, $0 \leq i \leq n$, s.t. $\mathcal{T}^* \not\models \bigsqcap \overline{\mathcal{E}_i} \sqsubseteq \neg C$ and $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_i} \sqsubseteq \neg(C \sqcap \neg D)$, hence $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_i} \sqcap C \sqsubseteq D$, i.e., $\langle \mathcal{T}, \mathcal{D} \rangle \vdash_r C \precsim D$. In the second case, we have $\mathcal{T}^* \models C \sqsubseteq \bot$, that implies $\langle \mathcal{T}, \mathcal{D} \rangle \vdash_r C \precsim D$.

From right to left: Assume $\langle \mathcal{T}, \mathcal{D}\rangle \vdash_r C \mathrel{\widetilde{\sqsubseteq}} D$. Then either there is an $i$ which is the lowest number s.t. $\mathcal{T}^* \not\models \bigsqcap \overline{\mathcal{E}_i} \sqsubseteq \neg C$ (hence $r(C) = i$), or $\mathcal{T}^* \models C \sqsubseteq \bot$. In the first case we have also that $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_i} \sqcap C \sqsubseteq D$, that implies that $\mathcal{T}^* \models \bigsqcap \overline{\mathcal{E}_i} \sqsubseteq \neg(C \sqcap \neg D)$, i.e. $r(C \sqcap \neg D) > i$. In the second case, $r(C) = \infty$, that implies $\langle \mathcal{T}, \mathcal{D}\rangle \models_R^{\cup} C \mathrel{\widetilde{\sqsubseteq}} D$. ■

## Appendix E. Proofs for Section 7

**Proposition 8** *Let $\mathcal{K} = \langle \mathcal{A}, \mathcal{T}, \mathcal{D}\rangle$ be a knowledge base, and $\mathcal{K}_{\mathcal{I}}$ the set of the individuals named in $\mathcal{A}$.*
*Given $\mathcal{K}$ and a linear order $s$ of the elements of $\mathcal{K}_{\mathcal{I}}$, Procedure RationalExtension determines a rational ABox extension of $\mathcal{K}$. Contrariwise, every rational ABox extension of $\mathcal{K}$ corresponds to the knowledge base generated by some linear order of the individuals in $\mathcal{K}_{\mathcal{I}}$.*

**Proof:**
The first part is immediate. For the second part, assume that there is a rational extension $\langle \mathcal{A}', \mathcal{T}\rangle$ of $\langle \mathcal{A}, \mathcal{T}, \mathcal{D}\rangle$ that cannot be generated by any sequence $s$ of the elements of $\mathcal{K}_{\mathcal{I}}$. $\mathcal{A}'$ associates to every individual $a \in \mathcal{K}_{\mathcal{I}}$ a concept from $\Delta = \langle \bigsqcap \overline{\mathcal{E}_0}, \dots, \bigsqcap \overline{\mathcal{E}_n}\rangle$, that we indicate as $\mathcal{E}^a$.

Let $\langle \mathcal{A}_{\mathcal{D}}, \mathcal{T}\rangle$ be a rational extension of $\langle \mathcal{A}, \mathcal{T}, \mathcal{D}\rangle$ that can be generated using a sequence of elements of $\mathcal{K}_{\mathcal{I}}$. The following procedure allows to define a sequence $s$ of the elements of $\mathcal{K}_{\mathcal{I}}$ s.t. $\langle \mathcal{A}_{\mathcal{D}}, \mathcal{T}\rangle$ can be generated using $s$, i.e., $\langle \mathcal{A}_{\mathcal{D}}, \mathcal{T}\rangle = \langle \mathcal{A}^s_{\mathcal{D}}, \mathcal{T}\rangle$.

Take each element of $\mathcal{K}_{\mathcal{I}}$ and associate with it the strongest default concept in $\Delta$ consistent with the knowledge base $\langle \mathcal{A}, \mathcal{T}\rangle$ (call it $\gamma^a$). Look for an individual $a \in \mathcal{K}_{\mathcal{I}}$ s.t. $\mathcal{E}^a = \gamma^a$, and consider $a$ the first element of the sequence $s$. Update $\mathcal{A}$ with $a : \mathcal{E}^a$, and repeat the procedure, until every individual has been associated to a $\Delta$-formula. With this procedure we can generate a sequence over the the individuals in $\mathcal{K}_{\mathcal{I}}$ that generates $\langle \mathcal{A}_{\mathcal{D}}, \mathcal{T}\rangle$ from $\langle \mathcal{A}, \mathcal{T}, \mathcal{D}\rangle$.

Since there is no sequence $s$ that can generate $\langle \mathcal{A}', \mathcal{T}\rangle$, the above procedure has to fail, that is, at some point it will not be possible to associate with any remaining individual $a$ a defeasible concept $\gamma^a$ s.t. $\mathcal{E}^a = \gamma^a$. That means that, for all the remaining $a$, $\mathcal{E}^a \neq \gamma^a$; for each such $a$, either $\models \mathcal{E}^a \sqsubseteq \gamma^a$ or $\models \gamma^a \sqsubseteq \mathcal{E}^a$. The first case is not possible, since $\langle \mathcal{A}', \mathcal{T}\rangle$ would be inconsistent ($\gamma^a$ has to be a maximally consistent default). Hence $\models \gamma^a \sqsubseteq \mathcal{E}^a$ and $\mathcal{E}^a \neq \gamma^a$ for all the remaining $a$. In such a case, $\langle \mathcal{A}', \mathcal{T}\rangle$ would not be a rational extension of $\langle \mathcal{A}_{\mathcal{D}}, \mathcal{T}\rangle$, since we could have another consistent model with at least the same amount of defeasible information associated to every individuals, and a larger amount associated to some of them. ■

**Proposition 9** *Given $\mathcal{K}$ and a linear order $s$ of the individuals in $\mathcal{K}$, the inference relation $\vdash^s_r$ satisfies the following properties:*

| | | |
|---|---|---|
| $(REF_{DL})$ | $\langle \mathcal{A}, \mathcal{T}, \Delta \rangle \vdash_r^s a : \tilde{C}$ for every $a : C \in \mathcal{A}$ | Reflexivity |
| $(LLE_{DL})$ | $\dfrac{\langle \mathcal{A} \cup \{b : D\}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C} \qquad \models D \equiv E}{\langle \mathcal{A} \cup \{b : E\}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C}}$ | Left Logical Equivalence |
| $(RW_{DL})$ | $\dfrac{\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C} \qquad \models C \sqsubseteq D}{\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{D}}$ | Right Weakening |
| $(CT_{DL})$ | $\dfrac{\langle \mathcal{A} \cup \{b : D\}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C} \qquad \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s b : \tilde{D}}{\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C}}$ | Cautious Transitivity (Cut) |
| $(CM_{DL})$ | $\dfrac{\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C} \qquad \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s b : \tilde{D}}{\langle \mathcal{A} \cup \{b : D\}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C}}$ | Cautious Monotonicity |
| $(OR_{DL})$ | $\dfrac{\langle \mathcal{A} \cup \{b : D\}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C} \qquad \langle \mathcal{A} \cup \{b : E\}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C}}{\langle \mathcal{A} \cup \{b : D \sqcup E\}, \mathcal{T}, \Delta \rangle \vdash_r^s a : \tilde{C}}$ | Left Disjunction |
| $(RM_{DL})$ | $\dfrac{\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C} \qquad \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \not\vdash_r^s b : \tilde{\neg D}}{\langle \mathcal{A} \cup \{b : D\}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C}}$ | Rational Monotonicity |

**Proof:**
For $REF_{DL}$, $LLE_{DL}$ and $RW_{DL}$ the proof is immediate. For $CT_{DL}$ and $CM_{DL}$, assume $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s b : \tilde{D}$, that is $\langle \mathcal{A}_{\mathcal{D}}^s, \mathcal{T} \rangle \models b : D$. Hence, for every $\bigsqcap \overline{\mathcal{E}_i} \in \Delta$ and every individual $z \in \mathcal{K}_{\mathcal{I}}$, $z : \bigsqcap \overline{\mathcal{E}_i}$ is consistent with $\langle \mathcal{A}, \mathcal{T} \rangle$ iff it is consistent with $\langle \mathcal{A} \cup \{b : D\}, \mathcal{T} \rangle$, and the procedure associates to each individual the same default formula either we start with $\mathcal{A}$ or with $\mathcal{A} \cup \{b : D\}$. So we have that $\langle \mathcal{A}_{\mathcal{D}}^s \cup \{b : D\}, \mathcal{T} \rangle = \langle (\mathcal{A} \cup \{b : D\})_{\mathcal{D}}^s, \mathcal{T} \rangle$ and $\langle \mathcal{A}_{\mathcal{D}}^s \cup \{b : D\}, \mathcal{T} \rangle \models a : C$ iff $\langle (\mathcal{A} \cup \{b : D\})_{\mathcal{D}}^s, \mathcal{T} \rangle \models a : C$. Since $\models$ satisfies $CT$ and $CM$, we have that $\langle \mathcal{A}_{\mathcal{D}}^s, \mathcal{T} \rangle \models a : C$ iff $\langle (\mathcal{A} \cup \{b : D\})_{\mathcal{D}}^s, \mathcal{T} \rangle \models a : C$, that is, $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C}$ iff $\langle \mathcal{A} \cup \{b : D\}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C}$.

For $OR_{DL}$, assume that $\langle \mathcal{A} \cup \{b : D\}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C}$, $\langle \mathcal{A} \cup \{b : E\}, \mathcal{T}, \mathcal{D} \rangle \vdash_r^s a : \tilde{C}$, and that $b$ is in the $n^{th}$ position in the sequence $s$. So, for the first $n-1$ elements of $s$ the association with the default-formulae is the same in both the models. For $b$, assume that the procedure assigns $b : \bigsqcap \overline{\mathcal{E}_i}$ in case $b : D$, and $cassb \bigsqcap \overline{\mathcal{E}_j}$ in case $b : E$. We can have $\bigsqcap \overline{\mathcal{E}_i} = \bigsqcap \overline{\mathcal{E}_i}$, $\models \bigsqcap \overline{\mathcal{E}_i} \sqsubseteq \bigsqcap \overline{\mathcal{E}_j}$, or $\models \bigsqcap \overline{\mathcal{E}_j} \sqsubseteq \bigsqcap \overline{\mathcal{E}_i}$. In the first case the procedure for the assignment of the defaults continues in the same way in both the knowledge bases, and it is the same also if we have $b : D \sqcup E$, that is, $\langle \mathcal{A} \cup \{b : D\}, \mathcal{T}, \mathcal{D} \rangle$, $\langle \mathcal{A} \cup \{b : E\}, \mathcal{T}, \mathcal{D} \rangle$, and $\langle \mathcal{A} \cup \{b : D \sqcup E\}, \mathcal{T}, \mathcal{D} \rangle$ are completed exactly with the same defeasible concepts, obtaining, respectively, the ABoxes $(\mathcal{A} \cup \{b : D\})_{\mathcal{D}}^s = \mathcal{A}' \cup \{b : D\}$, $(\mathcal{A} \cup \{b : E\})_{\mathcal{D}}^s, = \mathcal{A}' \cup \{b : E\}$, and $(\mathcal{A} \cup \{D \sqcup b : E\})_{\mathcal{D}}^s = \mathcal{A}' \cup \{D \sqcup b : E\}$, for some ABox $\mathcal{A}'$. So we have that $\mathcal{A}' \cup \{b : D\} \models a : C$ and $\mathcal{A}' \cup \{b : E\} \models a : C$, and, since $\models$ satisfies OR, we obtain $\mathcal{A}' \cup \{b : D \sqcup E\} \models a : C$, that is, $\langle (A \cup \{b : D \sqcup E\})_{\mathcal{D}}^s, \mathcal{T} \rangle \models a : C$. If $\models \bigsqcap \overline{\mathcal{E}_i} \sqsubseteq \bigsqcap \overline{\mathcal{E}_j}$ and $b : D \sqcup E$, the procedure associates to $b$ the strongest of the two defaults, that is, $\bigsqcap \overline{\mathcal{E}_i}$. Since $\bigsqcap \overline{\mathcal{E}_i}$ is not consistent with $b : E$, in every following consistency check the procedure will be forced to consider that $b : D$ holds, and the assignment of the defaults to the individuals will proceed as in the case where $b : D$ is in the KB, and $\langle \mathcal{A} \cup \{b : D \sqcup E\}, \mathcal{T}, \mathcal{D} \rangle$ will entail the same formulae as $\langle \mathcal{A} \cup \{b : D\}, \mathcal{T}, \mathcal{D} \rangle$. Analogously, if $\models \bigsqcap \overline{\mathcal{E}_j} \sqsubseteq \bigsqcap \overline{\mathcal{E}_i}$, the rational ABox extension of $\langle \mathcal{A} \cup \{b : D \sqcup E\}, \mathcal{T}, \mathcal{D} \rangle$ will correspond to the one of $\langle \mathcal{A} \cup \{b : E\}, \mathcal{T}, \mathcal{D} \rangle$.

Finally, for $RM_{DL}$, $b : D$ is consistent with $\langle \mathcal{A}_{\mathcal{D}}^s, \mathcal{T} \rangle$, so the presence of $b : D$ in the knowledge base does not influence the association of the defaults to the individuals, and

$\mathcal{A}^s_{\mathcal{D}} \subseteq (\mathcal{A} \cup \{b : D\})^s_{\mathcal{D}}$. Eventually, $\langle \mathcal{A}^s_{\mathcal{D}}, \mathcal{T} \rangle \models a : C$, that implies $\langle (\mathcal{A} \cup \{b : D\})^s_{\mathcal{D}}, \mathcal{T} \rangle \models a : C$, i.e. $\langle \mathcal{A} \cup \{b : D\}, \mathcal{T}, \mathcal{D} \rangle \vdash^s_r a : \tilde{C}$. ■

**Proposition 10** *Deciding $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \vdash^s_r a : \tilde{C}$ in $\mathcal{ALC}$ is an ExpTime-complete problem.*

**Proof:**
ABox decision in $\mathcal{ALC}$ is ExpTime-complete. $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$ is a knowledge base s.t. $\mathcal{D}$ is partitioned into $\mathcal{D}_0, \dots, \mathcal{D}_n$ and in the ABox are named $m$ individuals ($|\mathcal{K}_{\mathcal{I}}| = m$). Given a sequence $s$ of the individuals in $\mathcal{K}_{\mathcal{I}}$, to decide if $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \vdash^s_r a : \tilde{C}$ we need to do for each individual in $\mathcal{K}_{\mathcal{I}}$ at most $n$ ABox consistency checks to decide which default we can associate with that particular individual, and, eventually, once we have associated to each individual the strongest default possible, we have to check if $a : C$ is a classical consequence of the rational ABox extension. Hence, in the worst case we need $(n * m) + 1$ classical $\mathcal{ALC}$ decision steps, hence the procedure is ExpTime-complete. ■

**Proposition 12** *In the presence of a knowledge base $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$ that has a single rational ABox extension, checking the uniqueness of the rational ABox extension and, in case, whether $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle \vdash_r a : \tilde{C}$ is an ExpTime-complete problem in $\mathcal{ALC}$.*

**Proof:**
It is the same situation as in Proposition 10. We need at worst $n * m$ classical decision procedures to associate with each individual a default concept, then another one to check the overall consistency of the new knowledge base, and eventually, in case it is consistent, a last one to decide whether $a : C$ is a classical consequence of the rational ABox extension just defined. All in all, $(n * m) + 2$ ExpTime-complete decision procedures. ■

**Lemma 4** *A knowledge base $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$ in rank normal form is consistent if and only if $\langle \mathcal{A}, \mathcal{T} \rangle \not\models \top \sqsubseteq \bot$.*

**Proof:**

From left to right it is immediate. From right to left, we know from Corollary 2 that $\langle \mathcal{T}, \mathcal{D} \rangle$ has a ranked model iff $\mathcal{T} \not\models \top \sqsubseteq \bot$ (we have assumed that we have already moved all the infinitely ranked information possibly in $\mathcal{D}$ into $\mathcal{T}$, and $\mathcal{D}$ has been already ranked into $\mathcal{D}_0, \dots, \mathcal{D}_n$). Assume $\langle \mathcal{A}, \mathcal{T} \rangle \not\models \top \sqsubseteq \bot$; then, since $\mathcal{ALC}$satisfies the FMP (see Appendix G), there is a finite classical model $M = \langle \Delta^M, \cdot^M \rangle$ of $\langle \mathcal{A}, \mathcal{T} \rangle$. Let $R^{\mathcal{K}}_{\cup} = \langle \Delta^{R^{\mathcal{K}}_{\cup}}, \cdot^{R^{\mathcal{K}}_{\cup}}, \prec^{R^{\mathcal{K}}_{\cup}} \rangle$ be the characteristic model of the RC of $\langle \mathcal{T}, \mathcal{D} \rangle$ as defined in Section 5 (it must exist since $\langle \mathcal{T}, \mathcal{D} \rangle$ is consistent). Now, define a new interpretation $R^* = \langle \Delta^{R^*}, \cdot^{R^*}, \prec^{R^*} \rangle$ obtained merging the two interpretations in the following way:

- $\Delta^{R^*} = \Delta^{R^{\mathcal{K}}_{\cup}} \cup \Delta^M$;
- for every atomic concept $A$, $A^{R^*} = A^{R^{\mathcal{K}}_{\cup}} \cup A^M$;
- for every role $r$, $r^{R^*} = r^{R^{\mathcal{K}}_{\cup}} \cup r^M$;

- for every individual $a$ named in $\mathcal{A}$, $a^{R^*} = a^M$;
- for every object $o \in \Delta^M$, $h_{R^*}(o) = i$, with $i \leq n$, iff $M \Vdash o : \bigsqcap \overline{\mathcal{E}_i}$ and $M \not\Vdash o : \bigsqcap \overline{\mathcal{E}_{i-1}}$, or $h_{R^*}(o) = n+1$ if $M \not\Vdash o : \bigsqcap \overline{\mathcal{E}_n}$.

It is easy to check that $R'$ is a ranked model of $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$. ∎

**Lemma 5** Let $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$ be a consistent knowledge base in rank normal form, and $R^{\mathcal{K}}_{\cup}$ be the RC model of $\langle \mathcal{T}, \mathcal{D} \rangle$. Then there is at least an interpretation $R' = \langle \Delta^{R^{\mathcal{K}}_{\cup}}, \cdot^{R'}, \prec^{R^{\mathcal{K}}_{\cup}} \rangle$ s.t.:

- for every atomic concept $A$, $A^{R'} = A^{R^{\mathcal{K}}_{\cup}}$;
- for every role $r$, $r^{R'} = r^{R^{\mathcal{K}}_{\cup}}$;
- $R' \Vdash \mathcal{A}$.

**Proof:**

Consider the proof of Lemma 4. Form the FMP of the ranked models, form the existence of $R^*$ we know that there is also a finite model of $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$. From such a model, define a model of $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$ with a countably infinite $\Delta$ as a domain, and use it in the construction of $R^{\mathcal{K}}_{\cup}$.

∎

**Proposition 13** *Given a knowledge base $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$, each inference relation $\vdash^s_r$ defined by a sequence $s$ on the elements of $\mathcal{K}_{\mathcal{I}}$ corresponds to the entailment relation $\models^{\leq}_h$ for some $h$, and the other way around. The inference relation $\vdash_r$, corresponding to the intersection of all $\vdash^s_r$ generated by $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$, corresponds to the entailment relation $\models^{\leq}$.*

**Proof:**
Given a knowledge base $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$, in $R^{\mathcal{K}}_{\cup}$ there is a correspondence between the height of the objects and their ranking, that is, if an object $o$ has a height $i$, then the model associates to $o$ the defeasible information $\bigsqcap \overline{\mathcal{E}_i}$ (or nothing if $h_{R^{\mathcal{K}}_{\cup}}(o) = n+1$, with $n$ the highest value of $\bigsqcap \overline{\mathcal{E}_i}$). Hence, given all the minimal models of a knowledge base $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$ s.t. all the individuals in $\mathcal{K}_{\mathcal{I}}$ have the same height in each model, i.e., the models defining $\models^{\leq}_h$, we consider all the models that associate with each individual $x \in \mathcal{K}_{\mathcal{I}}$ a specific default concept $\delta_i$, s.t. it is not possible to associate a stronger default to each of them. This corresponds to the notion of rational ABox extension that, by Proposition 8, corresponds to the inference relation $\vdash^s_r$ generated by some sequence $s$. In the other direction, given a knowledge base $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$ and an inference relation $\vdash^s_r$, it corresponds to a rational ABox extension of $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$, and we can define the corresponding class of minimal models associating to each individual in $\mathcal{K}_{\mathcal{I}}$ the height $i$ if we have associated to it the defeasible information $\bigsqcap \overline{\mathcal{E}_i}$ (or the height $n+1$ if we have associated nothing to it).

The correspondence between $\vdash_r$ and $\models^{\leq}$ is an immediate consequence. ∎

## Appendix F. Proofs for Section 8

**Lemma 7** *Let $\mathscr{F}$ be the DTF defined by Procedure relaxSubsumption, and $\mathcal{O}$ a classical TBox, then $\mathscr{F}$ is a safe DTF for $\mathcal{O}$.*

**Proof:**

Suppose $\mathscr{D}$ is the special DTF which translates all subsumptions to defeasible ones. Suppose also that $C$ is totally exceptional w.r.t. $\mathscr{D}(\mathcal{O})$ for some $C$ but that $C$ is not totally exceptional w.r.t. $\mathscr{F}(\mathcal{O})$. We try to derive a contradiction from this.

From our supposition that $C$ is not totally exceptional w.r.t. $\mathscr{F}(\mathcal{O})$ we have two cases: either $C$ is not exceptional at all w.r.t. $\mathscr{F}(\mathcal{O})$ or $C$ is *normally* exceptional w.r.t. $\mathscr{F}(\mathcal{O})$.

*Case* 1*:* $C$ is not exceptional w.r.t. $\mathscr{F}(\mathcal{O})$. From our supposition that $C$ is totally exceptional w.r.t. $\mathscr{D}(\mathcal{O})$ we can infer that $\mathcal{O} \models C \sqsubseteq \bot$ from Lemma 6. Let $\mathcal{J}_1, \ldots, \mathcal{J}_n$ be the justifications for $\mathcal{O} \models C \sqsubseteq \bot$. Because we know that $C$ is totally exceptional w.r.t. $\mathscr{D}(\mathcal{O})$ it must be the case that for at least one $1 \leq i \leq n$, $C$ is totally exceptional w.r.t. $\mathscr{D}(\mathcal{J}_i)$. We can easily see from the depleting property of star locality based module that $\mathscr{D}(\mathcal{J}_i) \subseteq \mathscr{F}(\mathcal{O})$. Therefore $C$ is totally exceptional w.r.t. $\mathscr{F}(\mathcal{O})$. This is a contradiction and therefore it cannot be the case that $C$ is not exceptional w.r.t. $\mathscr{F}(\mathcal{O})$.

*Case* 2*:* $C$ is normally exceptional w.r.t. $\mathscr{F}(\mathcal{O})$. This is impossible because we have shown in Case 1 that there is a justification $\mathcal{J}_i$ for $\mathcal{O} \models C \sqsubseteq \bot$ s.t. $C$ is totally exceptional w.r.t. $\mathscr{D}(\mathcal{J}_i)$. Therefore $\mathscr{D}(\mathcal{J}_i) \subseteq \mathscr{F}(\mathcal{O})$ and it must be the case that $C$ is totally exceptional w.r.t. $\mathscr{F}(\mathcal{O})$.
■

## Appendix G. Finite Model Property for Defeasible $\mathcal{ALC}$ Knowledge Bases

We need to prove that defeasible $\mathcal{ALC}$ ontologies satisfy the Finite Model Property. To this end, consider a finite defeasible $\mathcal{ALC}$ knowledge base $\mathcal{K} = \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$, and let $R = \langle \Delta^R, \cdot^R, \prec^R \rangle$ be a ranked model of $\mathcal{K}$ (with $\Delta^R$ possibly infinite). Let $\preceq^R$ be the weak complete ordering generated from $\prec^R$, that is, $x \preceq^R y$ iff $y \not\prec^R x$. Let $\mathsf{N} \cup \mathsf{N}_{\mathscr{C}} \cup \mathsf{N}_{\mathscr{R}}$ be the signature of our language, and let $\Gamma$ be a set of concepts $\{C_1, \ldots, C_n\}$ s.t. $\Gamma$ is obtained closing under sub-concepts and negation the concepts appearing in the axioms in $\mathcal{K}$. Now we define the equivalence relation $\sim_\Gamma$ as

$$\forall x, y \in \Delta^R, x \sim_\Gamma y \text{ iff } \forall C \in \Gamma, x \in C^R \text{ iff } y \in C^R \ .$$

We indicate with $[x]_\Gamma$ the equivalence class of the individuals that are related to an individual $x$ through $\sim_\Gamma$:

$$[x]_\Gamma = \{y \in \Delta^R \mid x \sim_\Gamma y\} \ .$$

We introduce a new model $R' = \langle \Delta^{R'}, \cdot^{R'}, \preceq^{R'} \rangle$, defined as:

- $\Delta^{R'} = \{[x]_\Gamma \mid x \in \Delta^R\}$;
- for every $A \in \mathsf{N}_{\mathscr{C}} \cap \Gamma$, $A^{R'} = \{[x]_\Gamma \mid x \in A^R\}$;
- for every $A \notin \mathsf{N}_{\mathscr{C}} \cap \Gamma$, $A^{R'} = \emptyset$;
- for every $r \in \mathsf{N}_{\mathscr{R}}$, $r^{R'} = \{\langle [x]_\Gamma, [y]_\Gamma \rangle \mid \langle x, y \rangle \in r^R\}$;

- For every $[x]_\Gamma, [y]_\Gamma \in \Delta^{R'}$, $[x]_\Gamma \preceq^{R'} [y]_\Gamma$ if there is an object $z \in [x]_\Gamma$ s.t. for all the objects $v \in [y]_\Gamma$, $z \prec^R v$; otherwise, $[y]_\Gamma \preceq^{R'} [x]_\Gamma$.
- for every $a \in \mathsf{N}$, $a^{R'} = [x]_\Gamma$ iff $a^R = x$.

Let $\prec^{R'}$ and $\equiv^{R'}$ defined as usual:

- $[x]_\Gamma \prec^{R'} [y]_\Gamma$ iff $[x]_\Gamma \preceq^{R'} [y]_\Gamma$ and $[y]_\Gamma \not\preceq^{R'} [x]_\Gamma$;
- $[x]_\Gamma \equiv^{R'} [y]_\Gamma$ iff $[x]_\Gamma \preceq^{R'} [y]_\Gamma$ and $[y]_\Gamma \preceq^{R'} [x]_\Gamma$.

Note that $[x]_\Gamma \prec^{R'} [y]_\Gamma$ iff there is an object $z \in [x]_\Gamma$ s.t. for all the objects $v \in [y]_\Gamma$, $z \prec^R v$, while $[x]_\Gamma \equiv^{R'} [y]_\Gamma$ iff for every $z \in [x]_\Gamma$ there is a $v \in [y]_\Gamma$ s.t. $v \preceq^R z$, and for every $v \in [y]_\Gamma$ there is a $z \in [x]_\Gamma$ s.t. $z \preceq^R v$.

Given that $\Gamma$ is finite, $\Delta^{R'}$ is clearly finite. The following results are easy to prove.

**Lemma 32 ($*$)** *For every $C \in \Gamma$ and every $x \in \Delta^R$, $x \in C^R$ iff $[x]_\Gamma \in C^{R'}$.*

**Proof:**
The proof is straightforward by induction on the construction of the concepts. ■

**Lemma 33 ($*$)** *The preorder $\preceq^{R'}$ is a complete preorder.*

**Proof:**
Reflexivity: Assume $[x]_\Gamma \not\preceq^{R'} [x]_\Gamma$. By the definition of $\preceq^{R'}$, it implies that we could not have any $x' \in [x]_\Gamma$ s.t. $x' \prec^R x''$ for every $x'' \in [x]_\Gamma$, that in turn would imply $[x]_\Gamma \preceq^{R'} [x]_\Gamma$, and we would have an absurdity.

Transitivity: Assume $[x]_\Gamma \preceq^{R'} [y]_\Gamma$ and $[y]_\Gamma \preceq^{R'} [z]_\Gamma$. We have four possible cases.

1) Let $[x]_\Gamma \prec^{R'} [y]_\Gamma$ and $[y]_\Gamma \prec^{R'} [z]_\Gamma$. It follows that there is an individual $x' \in [x]_\Gamma$ s.t. $x' \prec^R y'$ for every $y' \in [y]_\Gamma$, and that there is a $y^* \in [y]_\Gamma$ s.t. $y^* \prec z'$ for every $z' \in [z]_\Gamma$. Hence $x' \prec^R z'$ for every $z' \in [z]_\Gamma$, and $[x]_\Gamma \prec^{R'} [z]_\Gamma$.

2) Assume that $[x]_\Gamma \prec^{R'} [y]_\Gamma$ and $[y]_\Gamma \equiv^{R'} [z]_\Gamma$. This implies that there is an individual $x' \in [x]_\Gamma$ s.t. $x' \prec^R y'$ for every $y' \in [y]_\Gamma$, that for every $y' \in [y]_\Gamma$ there is at least a $z' \in [z]_\Gamma$ s.t. $z' \preceq^R y'$, and that for every $z' \in [z]_\Gamma$ there is at least a $y' \in [y]_\Gamma$ s.t. $y' \preceq^R z'$. Then it must be the case that $x' \prec^R z'$ for every $z' \in [z]_\Gamma$, that is, $[x]_\Gamma \prec^{R'} [z]_\Gamma$; otherwise we would have that there is a $z' \in [z]_\Gamma$ s.t. $z' \preceq^R x'$, that, since there is a $y' \in [y]_\Gamma$ s.t. $y' \preceq^R z'$, by the transitivity of $\preceq^R$ would imply $y' \preceq^R x'$, against the hypothesis that $[x]_\Gamma \prec^{R'} [y]_\Gamma$.

3) Assume that $[x]_\Gamma \equiv^{R'} [y]_\Gamma$ and $[y]_\Gamma \prec^{R'} [z]_\Gamma$. There must be a $y' \in [y]_\Gamma$ s.t. $y' \prec^R z'$ for every $z' \in [z]_\Gamma$, and there must be an $x' \in [x]_\Gamma$ s.t. $x' \preceq^R y'$. Hence $x' \prec^R z'$ for every $z' \in [z]_\Gamma$, and we can conclude that $[x]_\Gamma \prec^{R'} [z]_\Gamma$.

4) Assume that $[x]_\Gamma \equiv^{R'} [y]_\Gamma$ and $[y]_\Gamma \equiv^{R'} [z]_\Gamma$. Then it is easy to check for every $x' \in [x]_\Gamma$ there must be a $z' \in [z]_\Gamma$ s.t. $z' \preceq x'$ and for every $z' \in [z]_\Gamma$ there must be an $x' \in [x]_\Gamma$ s.t. $x' \preceq c'$.

Completeness: It is sufficient to prove that for every pair of individuals $[x]_\Gamma, [y]_\Gamma$ in $R'$, either $[x]_\Gamma \preceq^{R'} [y]_\Gamma$ or $[y]_\Gamma \preceq^{R'} [x]_\Gamma$, which is an immediate consequence of the definition of $\preceq^{R'}$. ■

**Lemma 34 ($*$)** *The order $\prec^{R'}$ is a modular order that satisfies smoothness.*

**Proof:**
Irreflexivity and transitivity are immediate consequences of the properties of $\preceq^{R'}$ and the definition of $\prec^{R'}$ from $\preceq^{R'}$. Smoothness is an immediate consequence of the finiteness of the domain $\Delta^{R'}$.

We need to prove modularity, defining a ranking function over $\Delta^{R'}$ corresponding to the order $\prec^{R'}$. Due to the finiteness of the domain, the set $\min_{\prec^{R'}}(\Delta^{R'})$ is defined. Let $rk(x) = 0$ for every $x \in \min_{\prec^{R'}}(\Delta^{R'})$. Now, let $\Delta^{R'}_1 := \Delta^{R'} \setminus \min_{\prec^{R'}}(\Delta^{R'})$, and let $rk(x) = 1$ for every $x \in \min_{\prec^{R'}}(\Delta^{R'}_1)$ (again, $\min_{\prec^{R'}}(\Delta^{R'}_1)$ must be defined). Let $\Delta^{R'}_2 := \Delta^{R'} \setminus \min_{\prec^{R'}}(\Delta^{R'}_1)$ and $rk(x) = 2$ for every $x \in \min_{\prec^{R'}}(\Delta^{R'}_2)$, and so on, until we reach an $n$ s.t. $\Delta^{R'}_n = \emptyset$ (it must happen, since $\Delta^{R'}$ is finite). The function $rk$ is a ranking function $rk : X \longrightarrow \mathbb{Q}$ s.t. for every $x, y \in X$, $x \prec y$ iff $rk(x) < rk(y)$.

■

**Theorem 7 (∗) [FMP]** *Given a finite ontology $\mathcal{K} = \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$, if $\mathcal{K}$ has a ranked model $R$, then it has a* finite *ranked model $R'$.*

**Proof:**
Let $\mathcal{K} = \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$ be a defeasible ontology, $R$ a model of $\mathcal{K}$ and $R'$ a finite interpretation constructed from $R$ as defined above. The proof that $R'$ satisfies $\langle \mathcal{A}, \mathcal{T} \rangle$ is straightforward by Lemma 32 plus the observation that , for every $a, b \in \mathsf{N}$ and every $r \in \mathsf{N}_{\mathscr{R}}$, $R' \Vdash (a, b) : r$ iff $R \Vdash (a, b) : r$.

About $\mathcal{D}$, let $C \mathrel{\underset{\sim}{\sqsubset}} D \in \mathcal{D}$. Hence, either $C^R = \emptyset$, or the height of $C \sqcap D$ in $R$ is lower than the height of $C \sqcap \neg D$, that is, there is at least an individual $b$ satisfying $C \sqcap D$ s.t. for every individual $a$ satisfying $C \sqcap \neg D$, $b \prec^R a$. Since $C$, $D$ and $\neg D$ are in $\Gamma$, the individual $[y]_\Gamma \in \Delta^{R'}$ (obtained from $b \in \Delta^R$, that hence satisfies $C \sqcap D$) must be preferred to all the individuals satisfying $C \sqcap \neg D$, that is, $[y]_\Gamma \prec^{R'} [x]_\Gamma$ for every individual $[x]_\Gamma$ s.t. $[x]_\Gamma \in (C \sqcap \neg D)^{R'}$. Hence $R' \Vdash C \mathrel{\underset{\sim}{\sqsubset}} D$. ■

We can obtain also the analogous *finite counter-model property.*

**Proposition 15 (∗) [FCMP]** *Given a finite ontology $\mathcal{K} = \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$ and an inclusion axiom $C \mathrel{\underset{\sim}{\sqsubset}} D$, if $\mathcal{K}$ has a ranked model $R$ that is also a counter-model of $C \mathrel{\underset{\sim}{\sqsubset}} D$, then it has a* finite *ranked model $R'$ that is a counter-model of $C \mathrel{\underset{\sim}{\sqsubset}} D$.*

**Proof:**
It is sufficient to apply the same construction defined for the FMP: if $R$ is not a model of $C \mathrel{\underset{\sim}{\sqsubset}} D$, it means that there is an individual $a$ s.t. $R \Vdash a : C \sqcap \neg D$ and $a \preceq_R b$ for every individual $b$ s.t. $R \Vdash b : C \sqcap D$. That implies that in $R'$ $[x]_\Gamma \preceq_{R'} [y]_\Gamma$ for every $b$ s.t. $R \Vdash b : C \sqcap D$, and consequently $R' \not\Vdash C \mathrel{\underset{\sim}{\sqsubset}} D$. ■

**Corollary 5 (∗)** *Given a finite ontology $\mathcal{K} = \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$, if $\mathcal{K}$ has a ranked model $R$, then for every concept $C$ s.t. $h_R(C) = 0$ there is also a* finite *ranked model $R'$ of $\mathcal{K}$ s.t. $h_{R'}(C) = 0$.*

**Proof:**
Given $\mathcal{K} = \langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$ and a concept $C$ s.t. $h_R(C) = 0$, a finite model $R'$ satisfying the constraint above can be defined in the same way as the model $R'$ built here above to prove the FMP; we just need to add $C$ to the set $\Gamma$ (and close $\Gamma$ also under the subconcepts of $C$

and their negations). To see that $R'$ is a model of $\mathcal{K}$ just go again through the FMP proof above, and check that the addition of $C$ in $\Gamma$ does not affect any of the above proofs.

$h_R(C) = 0$ implies that there is an object $x \in \Delta^R$ s.t. $x \in C^R$ and $h_R(x) = 0$. Now consider $[x]_\Gamma$; by Lemma 32, $[x]_\Gamma \in C^{R'}$. Since $h_R(x) = 0$, for every $[y]_\Gamma \in \Delta^{R'}$ it cannot be the case that there is an object $z \in [y]_\Gamma$ s.t. $z \prec^R v$ for every $v \in [x]_\Gamma$; hence, the definition of $\preceq^{R'}$ implies that for every $[y]_\Gamma \in \Delta^{R'}$, $[x]_\Gamma \preceq^{R'} [y]_\Gamma$, that is, $h_{R'}([x]_\Gamma) = 0$, that implies $h_{R'}(C) = 0$. ■