On the Logical Properties of the Nonmonotonic Description Logic $\mathcal{DL}^{\rm N}$

P. A. Bonatti, L. Sauro
Dep. of Electrical Engineering and Information Technologies
Università di Napoli Federico II
{pab,luigi.sauro}@unina.it

April 25, 2018

Abstract

$\mathcal{DL}^{\rm N}$ is a recent nonmonotonic description logic, designed for satisfying independently proposed knowledge engineering requirements, and for removing some recurrent drawbacks of traditional nonmonotonic semantics. In this paper we study the logical properties of $\mathcal{DL}^{\rm N}$ and illustrate some of the relationships between the KLM postulates and the characteristic features of $\mathcal{DL}^{\rm N}$, including its novel way of dealing with unresolved conflicts between defeasible axioms. Moreover, we fix a problem affecting the original semantics of $\mathcal{DL}^{\rm N}$ and accordingly adapt the reduction from $\mathcal{DL}^{\rm N}$ inferences to classical inferences. Along the paper, we use various versions of the KLM postulates to deepen the comparison with related work, and illustrate the different tradeoffs between opposite requirements adopted by each approach.

1 Introduction

Recently, in [5], a new nonmonotonic description logic called $\mathcal{DL}^{\rm N}$ has been introduced with the goal of supporting ontology authoring by means of nonmonotonic reasoning. $\mathcal{DL}^{\rm N}$ aims at removing some recurrent drawbacks of traditional nonmonotonic DLs[1] such as: (i) *inheritance blocking*, that is a drawback of preferential semantics and rational closure; (ii) undesired *closed world assumption effects*, that affect circumscription, typicality logic and some probabilistic logics; (iii) the inability to specify whether roles shall range over normal/prototypical individuals or not, that affects most nonmonotonic DLs. Moreover, $\mathcal{DL}^{\rm N}$ adopts a novel conflict resolution mechanism that helps in detecting unresolved conflicts between mutually inconsistent defaults. Since unresolved conflicts frequently correspond to missing knowledge, highlighting such conflicts constitutes an important support to knowledge base debugging and validation. A further useful property of $\mathcal{DL}^{\rm N}$ is that it can be translated into classical DLs, so that its implementations can rely on mature and well-optimized inference engines.

[1] Abbreviation for "description logics".

The translation can be computed in polynomial time and does not involve complex constructs, so $\mathcal{DL}^{\mathrm{N}}$ preserves the tractability of the two major low-complexity families of DLs, that is $\mathcal{EL}$ and *DL-lite*, that correspond to the EL and QL profiles of OWL2, respectively. The experiments in [5, 11] show unparalleled scalability properties over large knowledge bases, with more than 10^5 axioms. The relationships between the above features and the knowledge engineering requirements that have been independently introduced in the fields of biomedical ontologies and declarative policy languages have been extensively discussed in [5].

The number and complexity of the issues dealt with in [5] led us to postpone the analysis of the logical properties of $\mathcal{DL}^{\mathrm{N}}$. The main goal of this paper is precisely studying the logical properties of its consequence relation, defeasible axioms and normal instances. Our analysis includes a comparison of $\mathcal{DL}^{\mathrm{N}}$'s inferences with verbatim and internalized versions of the KLM postulates [29]. For our purposes, these postulates are not necessarily desiderata, due to the loose correspondence between their motivations and $\mathcal{DL}^{\mathrm{N}}$'s goals and semantics. However, we regard them as a useful technical tool for profiling and comparing the behavior of different logics, since the validity (or non-validity) of the postulates has been extensively investigated in the context of most nonmonotonic logics.

As a second contribution, we fix a problem affecting the original version of $\mathcal{DL}^{\mathrm{N}}$, that does not fully satisfy an internalized version of the KLM postulate LLE (*left logical equivalence*) and fails to derive some natural equivalences. The refined semantics improves also another foundational aspect: while the original definition of overriding quantifies over the class of all interpretations, the refined version depends only on a proper *set* of interpretations, thereby preventing any form of *unrestricted comprehension* [27]. We adapt the translation of $\mathcal{DL}^{\mathrm{N}}$ into classical description logic accordingly and prove its correctness, fixing a minor error in [5].

Finally, we leverage on two features that make $\mathcal{DL}^{\mathrm{N}}$ an excellent playground for analysing the interplay of the KLM postulates with other principles and requirements. First, $\mathcal{DL}^{\mathrm{N}}$'s syntax is relatively rich: it is a first-order language, where standard individuals can be denoted through suitable concept expressions. This feature makes it possible to encode examples that cannot be modelled in several other nonmonotonic DLs (nor in the original, propositional KLM framework). The second useful feature is that the underlying monotonic semantics is relatively weak, in comparison with other richer formalisms such as typicality logic; all logical properties follow from classical logic and the elementary principle that *all defaults should be applied unless they are contradicted by a group of higher priority (possibly non-defeasible) axioms*. This helps in identifying what needs to be added to the aforementioned principle – that is the core of default reasoning [37, 30] – if one desired to satisfy the KLM postulates in a default description logic. Using these technical tools, we shall compare $\mathcal{DL}^{\mathrm{N}}$ with other similar DLs in terms of the different tradeoffs between opposite desiderata stemming from the postulates, the above principle, and the emerging knowledge engineering requirements mentioned above.

In the next section, we recall some preliminaries on description logics, $\mathcal{DL}^{\mathrm{N}}$, and the KLM postulates, plus the basics of the two families of nonmonotonic DLs, that will be used as term of comparison in discussing logical properties and different ways of achieving them. In Section 3, we study the properties of the consequence relation of

Name	Syntax	Semantics
inverse role	R^-	$\{(d,e) \mid (e,d) \in R^{\mathcal{I}}\}$
universal role	U	$\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$
top	$\top$	$\Delta^{\mathcal{I}}$
bottom	$\bot$	$\emptyset$
nominal	$\{a\}$	$\{a^{\mathcal{I}}\}$
negation	$\neg C$	$\Delta^{\mathcal{I}} \setminus C^{\mathcal{I}}$
conjunction	$C \sqcap D$	$C^{\mathcal{I}} \cap D^{\mathcal{I}}$
disjunction	$C \sqcup D$	$C^{\mathcal{I}} \cup D^{\mathcal{I}}$
$\exists$ restriction	$\exists R.C$	$\{d \in \Delta^{\mathcal{I}} \mid \exists e \in \Delta^{\mathcal{I}}.[(d,e) \in R^{\mathcal{I}} \wedge e \in \mathcal{C}^{\mathcal{I}}]\}$
$\forall$ restriction	$\forall R.C$	$\{d \in \Delta^{\mathcal{I}} \mid \forall e \in \Delta^{\mathcal{I}}.[(d,e) \in R^{\mathcal{I}} \rightarrow e \in \mathcal{C}^{\mathcal{I}}]\}$

Figure 1: Syntax and semantics of the constructs used in the examples.

$\mathcal{DL}^{\mathrm{N}}$. Then a few sections analyze the logical behavior of expressions that are syntactically similar to constructs of other logics. In particular, Section 4 illustrates the behavior of defeasible inclusions by analogy with the defaults by Reiter and Lehmann. Section 5 describes the properties of the concepts that denote standard instances, and discusses internalized KLM postulates in relation to $\mathcal{DL}^{\mathrm{N}}$'s characteristic features. In Section 6 we refine the semantics of $\mathcal{DL}^{\mathrm{N}}$ to remove the aforementioned drawbacks. Finally, Section 7 completes the review of related work and Section 8 summarizes the results of the paper. The reader is assumed to be moderately familiar with nonmonotonic logics and description logics. Proofs have been moved to the appendix to improve readability.

2 Preliminaries

2.1 Description Logics

In DLs, *concepts* are inductively defined with the help of a set of *constructors*, starting with a set $\mathsf{N_C}$ of *concept names*, a set $\mathsf{N_R}$ of *role names*, and a set $\mathsf{N_I}$ of *individual names* (all countably infinite). We use the term *predicate* to refer to elements of $\mathsf{N_C} \cup \mathsf{N_R}$. Metavariables A, B will range over concept names, C and D over (possibly compound) concepts, R and S over roles, and a, b and d over individual names.

Expressive DLs support a large number of constructors; Figure 1 illustrates some of them that will be used throughout the examples of this paper.[2] Additionally, we shall adopt the usual abbreviation: $\exists R = \exists R.\top$.

The semantics of DLs is defined in terms of *interpretations* $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$. The *domain* $\Delta^{\mathcal{I}}$ is a non-empty set of individuals and the *interpretation function* $\cdot^{\mathcal{I}}$ maps each concept name $A \in \mathsf{N_C}$ to a subset $A^{\mathcal{I}}$ of $\Delta^{\mathcal{I}}$, each role name $R \in \mathsf{N_R}$ to a binary

[2]Note, however, that our framework applies also to richer DLs supporting fixpoint operators, full number restrictions, complex role inclusions, and all of the other operators not occurring in Figure 1.

relation $R^{\mathcal{I}}$ on $\Delta^{\mathcal{I}}$, and each individual name $a \in \mathsf{N_I}$ to an individual $a^{\mathcal{I}} \in \Delta^{\mathcal{I}}$. The extension of $\cdot^{\mathcal{I}}$ to inverse roles and some common compound concepts is inductively defined as shown in the third column of Figure 1. An interpretation $\mathcal{I}$ is called a *model* of a concept C if $C^{\mathcal{I}} \neq \emptyset$. If $\mathcal{I}$ is a model of C, we also say that C is *satisfied* by $\mathcal{I}$.

A *(general) TBox* is a finite set of *concept inclusions (CIs)* $C \sqsubseteq D$. As usual, we use $C \equiv D$ as an abbreviation for $C \sqsubseteq D$ and $D \sqsubseteq C$. An *ABox* is a finite set of *concept assertions* $C(a)$ and *role assertions* $R(a, b)$. An interpretation $\mathcal{I}$ *satisfies* (i) a concept inclusion $C \sqsubseteq D$ if $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$, (ii) an assertion $C(a)$ if $a^{\mathcal{I}} \in C^{\mathcal{I}}$, and (iii) an assertion $R(a, b)$ if $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in R^{\mathcal{I}}$. Then, $\mathcal{I}$ is a (classical) *model* of a TBox $\mathcal{T}$ (resp. an ABox $\mathcal{A}$) if $\mathcal{I}$ satisfies all the members of $\mathcal{T}$ (resp. $\mathcal{A}$).

In this paper we mention some important DLs that have been extensively studied in the literature and constitute the foundation of semantic web standards. The logic $\mathcal{ALC}$ is defined by the following grammar, where R ranges over role names:

$$C, D ::= A \mid \top \mid \bot \mid \neg C \mid C \sqcap D \mid C \sqcup D \mid \exists R.C \mid \forall R.C\,,$$

The logic $\mathcal{EL}$ supports only $\top$, $\sqcap$, and $\exists$. Its extension $\mathcal{EL}^{\bot}$ supports also $\bot$. The logic $\mathcal{EL}^{++}$ further adds *concrete domains* and some expressive role inclusions (see [2] for further details).

The logic DL-lite [18] supports inclusions shaped like $C \sqsubseteq D$ and $C \sqsubseteq \neg D$, where C and D range over concept names and *unqualified existential restrictions* such as $\exists R$ and $\exists R^-$ (where $R \in \mathsf{N_R}$). $\mathcal{EL}^{++}$ and DL-lite, respectively, constitute the foundation of the OWL2 profiles OWL2-EL and OWL2-QL. Both play an important role in applications; their inference problems are tractable (the same holds for some extensions of DL-lite, see [1]).

Finally we mention the very expressive DL $\mathcal{SROIQ}$ that constitutes the foundation of the full standard OWL2. Inference in this logic is 2NExpTime-complete [28].

$\mathcal{SROIQ}$ supports the universal role U that, in turn, can express boolean combinations of inclusions and assertions. For example, $\neg(C \sqsubseteq D)$ can be expressed as $\top \sqsubseteq \exists U.(C \sqcap \neg D)$, and $(C_1 \sqsubseteq D_1) \vee (C_2 \sqsubseteq D_2)$ can be expressed as $\top \sqsubseteq \big(\forall U.(\neg C_1 \sqcup D_1)\big) \sqcup \big(\forall U.(\neg C_2 \sqcup D_2)\big)$.

Some boolean combinations can be expressed without U. For instance, with the help of an auxiliary role, $\neg(C \sqsubseteq D)$ can be expressed as $\top \sqsubseteq \exists aux.(C \sqcap \neg D)$. Moreover, $(\{a\} \sqsubseteq \{b\}) \vee (\{a\} \sqsubseteq \{c\})$ is equivalent to $\{a\} \sqsubseteq \{b, c\}$ and

$$\begin{aligned} \neg(C(a)) \quad &\Leftrightarrow \quad (\neg C)(a)\,, \\ C(a) \vee D(a) \quad &\Leftrightarrow \quad (C \sqcup D)(a)\,. \end{aligned}$$

Some boolean combinations of DL axioms occur in the KLM postulates and play a role in some of our results.

2.2 The basics of $\mathcal{DL}^{\mathrm{N}}$

Let $\mathcal{DL}$ be any classical description logic language and let $\mathcal{DL}^{\mathrm{N}}$ be the extension of $\mathcal{DL}$ with a new concept name $\mathrm{N}C$ for each $\mathcal{DL}$ concept C. The new concepts are called *normality concepts* and denote the standard instances of C.

A *canonical* $\mathcal{DL}^{\rm N}$ *knowledge base* (hereafter *knowledge base*, for simplicity) is a disjoint union $\mathcal{KB} = \mathcal{S} \cup \mathcal{D}$ where $\mathcal{S}$ is a finite set of $\mathcal{DL}$ inclusions and assertions (called *strong* or classical axioms) and $\mathcal{D}$ is a finite set of *defeasible inclusions* (DIs, for short) that are expressions $C \sqsubseteq_n D$ where C is a $\mathcal{DL}$ concept and D a $\mathcal{DL}^{\rm N}$ concept.[3] If $\delta = (C \sqsubseteq_n D)$, then $\mathsf{pre}(\delta)$ and $\mathsf{con}(\delta)$ denote C and D, respectively.

The informal meaning of $C \sqsubseteq_n D$, roughly speaking, is: "*by default, standard instances satisfy* $C \sqsubseteq D$*, unless stated otherwise*", that is, unless some higher priority (possibly strong) axioms entail that standard instances belong to $C \sqcap \neg D$; in that case, $C \sqsubseteq_n D$ is *overridden*. The standard instances of a concept C are required to satisfy all the DIs that are not overridden in C. Accordingly, we will call the set of DIs satisfied by $\mathrm{N}C$ the *prototype* associated to C.

DIs are prioritized by a strict partial order $\prec$. If $\delta_1 \prec \delta_2$, then δ_1 has higher priority than δ_2. $\mathcal{DL}^{\rm N}$ solves automatically only the conflicts that can be settled using $\prec$. Any other conflict shall be resolved by the knowledge engineer (typically by adding specific DIs). Two priority relations have been investigated so far. Both are based on *specificity*: the specific default properties of a concept C have higher priority than the more generic properties of its superconcepts (i.e. those that subsume C). The priority relation used in most of the examples of [5] collects those superconcepts with strong axioms only:

$$\delta_1 \prec \delta_2 \text{ iff } \mathsf{pre}(\delta_1) \sqsubseteq_{\mathcal{S}} \mathsf{pre}(\delta_2) \text{ and } \mathsf{pre}(\delta_2) \not\sqsubseteq_{\mathcal{S}} \mathsf{pre}(\delta_1)\,.^4 \tag{1}$$

The second priority relation investigated in [5] is

$$\delta_1 \prec \delta_2 \text{ iff } \mathit{rank}(\delta_1) > \mathit{rank}(\delta_2), \tag{2}$$

where $\mathit{rank}(\cdot)$ is shown in Algorithm 1 and corresponds to the ranking function of rational closure adopted in [20, 22]. This relation uses also DIs to determine superconcepts, so (roughly speaking) a DI $C \sqsubseteq_n D$ – besides defining a default property for C – gives the specific default properties of C higher priority than those of D. The advantage of this priority relation is that it resolves more conflicts than (1); the main advantage of (1) is predictability; e.g. the effects of adding default properties to an existing, classical KB are more predictable, as the hierarchy used for determining specificity and resolving conflicts is the original, validated one, and is not affected by the new DIs (see also the related discussion in [6, 7], that adopt (1)).

The expression $\mathcal{KB} \mathrel{|\!\!\approx} \alpha$ means that α is a $\mathcal{DL}^{\rm N}$ *consequence* of $\mathcal{KB}$. Here we do not report the model-theoretic definition of $\mathrel{|\!\!\approx}$ (an improved version will be given in Sec. 6) and present only its reduction to classical reasoning. For all $\mathcal{DL}^{\rm N}$ subsumptions and assertions α, $\mathcal{KB} \mathrel{|\!\!\approx} \alpha$ holds iff $\mathcal{KB}^{\Sigma} \models \alpha$, where Σ is any set of normality concepts that contains at least the normality concepts that explicitly occur in $\mathcal{KB} \cup \{\alpha\}$, and $\mathcal{KB}^{\Sigma}$ is a classical knowledge base obtained as follows (recall that $\mathcal{KB} = \mathcal{S} \cup \mathcal{D}$):

First, for all DIs $\delta \in \mathcal{D}$ and all $\mathrm{N}C \in \Sigma$, let:

$$\delta^{\mathrm{N}C} = \big(\mathrm{N}C \sqcap \mathsf{pre}(\delta) \sqsubseteq \mathsf{con}(\delta)\big)\,. \tag{3}$$

---

[3] The definitions and results in [5] cover a wider class of knowledge bases, that here we call *general*, where normality concepts are allowed to occur in $\mathcal{S}$. The unrestricted use of normality concepts in strong axioms may override their intended semantics, that should be primarily driven by defeasible inclusions, as discussed in the note on usage below. A careful use of general knowledge bases, though, can be useful to tune the logical properties of $\mathcal{DL}^{\rm N}$, as discussed later in remarks 1 and 2.

[4] As usual, $C \sqsubseteq_{\mathcal{S}} D$ means that $\mathcal{S} \models C \sqsubseteq D$.

Algorithm 1: Ranking function

Input: Ontology $\mathcal{KB} = \mathcal{S} \cup \mathcal{D}$
Output: the function $\mathit{rank}(\cdot)$

$i := -1$; $\mathcal{E}_0 := \{C \sqsubseteq D \mid C \sqsubseteq_n D \in \mathcal{D}\}$;
repeat
  $i := i + 1$;
  $\mathcal{E}_{i+1} := \{C \sqsubseteq D \in \mathcal{E}_i \mid \mathcal{S} \cup \mathcal{E}_i \models C \sqsubseteq \bot\}$;
  forall $C \sqsubseteq_n D$ *s.t.* $C \sqsubseteq D \in \mathcal{E}_i \setminus \mathcal{E}_{i+1}$ **do**
    assign $\mathit{rank}(C \sqsubseteq_n D) := i$;
until $\mathcal{E}_{i+1} = \mathcal{E}_i$;
forall $C \sqsubseteq_n D \in \mathcal{E}_{i+1}$ **do** assign $\mathit{rank}(C \sqsubseteq_n D) := \infty$;
return $\mathit{rank}(\cdot)$;

The informal meaning of $\delta^{\mathrm{N}C}$ is: "$\mathrm{N}C$'s instances satisfy δ".

Second, let $\downarrow_{\prec\delta}$ be the operator that removes from axiom sets all the $\delta_0^{\mathrm{N}C}$ that do not have higher priority than δ, protecting the strong axioms of $\mathcal{KB}$:[5]

$$\text{for all } \mathcal{S}' \supseteq \mathcal{S}, \quad \mathcal{S}' \downarrow_{\prec\delta} = \mathcal{S} \cup (\mathcal{S}' \setminus \{\delta_0^{\mathrm{N}C} \mid \mathrm{N}C \in \Sigma \wedge \delta_0 \not\prec \delta\}) .$$

Third, let $\delta_1, \ldots, \delta_{|\mathcal{D}|}$ be any *linearization* of $(\mathcal{D}, \prec)$.[6]

Finally, let $\mathcal{KB}^{\Sigma} = \mathcal{KB}^{\Sigma}_{|\mathcal{D}|}$, where the sequence $\mathcal{KB}^{\Sigma}_i$ ($i = 1, 2, \ldots, |\mathcal{D}|$) is inductively defined as follows:

$$\mathcal{KB}^{\Sigma}_0 = \mathcal{S} \cup \big\{\mathrm{N}C \sqsubseteq C \mid \mathrm{N}C \in \Sigma\big\} \tag{4}$$

$$\mathcal{KB}^{\Sigma}_i = \mathcal{KB}^{\Sigma}_{i-1} \cup \big\{\delta_i^{\mathrm{N}C} \mid \mathrm{N}C \in \Sigma, \text{ and } \mathcal{KB}^{\Sigma}_{i-1} \downarrow_{\prec\delta_i} \cup \{\delta_i^{\mathrm{N}C}\} \not\models \mathrm{N}C \sqsubseteq \bot\big\} \tag{5}$$

(note that $\mathcal{KB}_i \subseteq \mathcal{KB}_{i+1}$, $i = 1, 2, \ldots, |\mathcal{D}|$).

In other words, the above sequence starts with $\mathcal{KB}$'s strong axioms extended with the inclusions $\mathrm{N}C \sqsubseteq C$, then processes the DIs δ_i in non-increasing priority order. If δ_i can be consistently added to C's prototype, given all higher priority DIs selected so far (which is verified by checking that $\mathrm{N}C \not\sqsubseteq \bot$ in line (5)), then its translation $\delta_i^{\mathrm{N}C}$ is included in $\mathcal{KB}^{\Sigma}$ (i.e. δ_i enters C's prototype), otherwise δ_i is discarded, and we say that δ_i is *overridden in* $\mathrm{N}C$.

We conclude this section with two technical remarks. First, in order to answer a given query α, it suffices to use the $\mathcal{KB}^{\Sigma}$ obtained from the smallest possible Σ, that contains only the normality concepts explicitly occurring in $\mathcal{KB}$ and α; however, larger Σ may be optionally used to pre-compute the properties of frequently used normality concepts, for optimization purposes (i.e. reusing the same $\mathcal{KB}^{\Sigma}$ for many queries). Second, choosing the linearization is not an issue, since they all yield the same result (cf. [5, Remark 2]).

[5] This is essential for the correctness of $\mathcal{KB}^{\Sigma}$ and corrects a minor error in [5].

[6] That is, $\{\delta_1, \ldots, \delta_{|\mathcal{D}|}\} = \mathcal{D}$ and for all $i, j = 1, \ldots, |\mathcal{D}|$, if $\delta_i \prec \delta_j$ then $i < j$.

A note on usage

$\mathcal{DL}^{\mathrm{N}}$ has two constructs for handling nonmonotonic inferences: defeasible inclusions and normality concepts. Defeasible inclusions are meant to specify the default properties of the standard instances denoted by normality concepts; so DIs do *not* affect the other predicates. For example, even if standard birds fly by default, still concept Bird contains instances that do not fly. Concepts Bird and $\mathrm{N}\mathit{Bird}$ are not induced to be as similar as possible by $\mathcal{DL}^{\mathrm{N}}$'s semantics, in order to avoid undesired closed-world effects, cf. [5].[7]

Conversely, a normality concept $\mathrm{N}C$ is meant to be axiomatized by the strong (monotonic) properties of C plus the default properties specified by DIs. So the correct way of asserting that *standard birds fly by default* is by means of the defeasible inclusion $\mathrm{N}\mathit{Bird} \sqsubseteq_n \mathit{Flying}$ and *not* through the strong inclusion $\mathrm{N}\mathit{Bird} \sqsubseteq \mathit{Flying}$, as the latter cannot be possibly overridden. This explains the rationale behind canonical knowledge bases, that by definition prevent the usage of normality concepts in strong axioms.

Since the effects of defeasible inclusions are confined to normality concepts, if a query Q contains no such concepts, then it is about certain, valid knowledge; accordingly, Q is entailed by $\mathcal{KB}$ if and only it follows classically from the strong axioms of $\mathcal{KB}$. On the contrary, if we are interested in inspecting the defeasible consequences of $\mathcal{KB}$, then Q should contain some normality concepts. For instance, in order to check whether standard birds fly by default, Q should be $\mathrm{N}\mathit{Bird} \sqsubseteq \mathit{Flying}$. Here $\mathit{Bird} \sqsubseteq_n \mathit{Flying}$ is *not* the correct query to pose, because even if a DI is a consequence of $\mathcal{KB}$, it can be *overridden*, and in this case it has no effect (a DI δ is entailed if for all normality concepts $\mathrm{N}C$, δ is either satisfied by the instances $\mathrm{N}C$ or overridden in $\mathrm{N}C$, cf. [5] and Section 6). This behavior of DIs is further discussed in Section 4. Inferring DIs may turn out to be useful for other purposes related to optimization, such as removing redundant DIs from a knowledge base.

Summarizing, in the typical examples of nonmonotonic reasoning occurring in the literature, DIs occur only in the knowledge base, while strong inclusions involving normality concepts occur only in queries.

Finally, a caveat on the intended meaning of normality concepts. Their name suggests that they are meant to collect the typical or normal instances of a concept, while their intended usage is actually more general, which is why we rather use the term "standard instances". To illustrate this idea, we anticipate one of the examples below, namely, the default inheritance of drug contraindications. It has been proposed by Rector not really because those contraindications normally hold, but rather because it is less dangerous to infer more contraindications than missing some of them. In this case, the properties of normality concepts are not determined by what we expect to hold in the world, but rather by what we deem safer to assume. Another opportunistic approach like this – where default properties are determined by what is more useful to assume – will be illustrated in Example 7. This category of intended applications

[7] For example, Circumscription would minimize the concept Whale because whales are abnormal mammals that live in the sea. Then Whale would contain only the individuals explicitly asserted to be whales. In the absence of such individuals, $\mathit{Whale} \sqsubseteq \bot$ holds. If we only assert $\mathit{Whale}(\mathit{Moby})$, then $\mathit{Whale} \equiv \{\mathit{Moby}\}$ holds.

departs from the goal of modelling what normally holds, that is the major motivation behind the KLM postulates.

Examples

We start by illustrating $\mathcal{DL}^{\mathrm{N}}$'s conflict handling approach, that has been motivated by knowledge engineering requirements stemming from biomedical applications and from the needs of semantic web policies. Most other logics silently neutralize the conflicts between nonmonotonic axioms with equal or incomparable priorities by computing the inferences that are invariant across all possible ways of resolving the conflict. A knowledge engineer, using her domain knowledge, may want to solve the conflict in favor of *some* of its possible resolutions, by adding specific axioms; however, if the logic silently neutralizes the conflict, then such missing knowledge may remain undetected and unfixed. This approach may cause serious problems in the policy domain:

Example 1 Suppose that project coordinators are both administrative staff and research staff. By default, administrative staff are allowed to sign payments, while research staff are not. A conflict arises since both of these default policies apply to project coordinators. Formally, $\mathcal{KB}$ can be formalized with:

$$\mathtt{Admin} \sqsubseteq_{\mathsf{n}} \exists \mathtt{has\_right.Sign} \quad (6)$$
$$\mathtt{Research} \sqsubseteq_{\mathsf{n}} \neg\exists \mathtt{has\_right.Sign} \quad (7)$$
$$\mathtt{PrjCrd} \sqsubseteq \mathtt{Admin} \sqcap \mathtt{Research} \quad (8)$$

Leaving the conflict unresolved may cause a variety of security problems. If project coordinators should *not* sign payments, and the default policy is *open* (authorizations are granted by default), then failing to infer $\neg\exists\mathtt{has\_right.Sign}$ would improperly authorize signatures. Conversely, if the authorization is to be granted and the policy is closed, then failing to prove $\exists\mathtt{has\_right.Sign}$ causes a *denial of service* (the user is unable to complete a legal operation). To prevent these problems, $\mathcal{DL}^{\mathrm{N}}$ makes the conflict visible by inferring $\mathcal{KB} \mathrel{|\!\!\approx} \mathrm{N}\,\mathtt{PrjCrd} \sqsubseteq \bot$ (i.e. making $\mathtt{PrjCrd}$'s prototype inconsistent). Technically, this can be proved by checking that $\mathcal{KB}^{\Sigma} \models \mathrm{N}\,\mathtt{PrjCrd} \sqsubseteq \bot$, where $\Sigma = \{\mathrm{N}\,\mathtt{PrjCrd}\}$. Here $\mathcal{KB}^{\Sigma}$ consists of (8), $\mathrm{N}\,\mathtt{PrjCrd} \sqsubseteq \mathtt{PrjCrd}$, and the following translations of (6) *and* (7) (none of which overrides the other because none is more specific under any of the two priorities):

$$\mathrm{N}\,\mathtt{PrjCrd} \sqcap \mathtt{Admin} \sqsubseteq \exists\mathtt{has\_right.Sign},$$
$$\mathrm{N}\,\mathtt{PrjCrd} \sqcap \mathtt{Research} \sqsubseteq \neg\exists\mathtt{has\_right.Sign}.$$

Now it is easy to see that $\mathrm{N}\,\mathtt{PrjCrd}$ is indeed inconsistent in $\mathcal{KB}^{\Sigma}$. Given this warning, a knowledge engineer can easily add the missing information by asserting either $\mathtt{PrjCrd} \sqsubseteq_{\mathsf{n}} \exists\mathtt{has\_right.Sign}$ or $\mathtt{PrjCrd} \sqsubseteq_{\mathsf{n}} \neg\exists\mathtt{has\_right.Sign}$. ∎

The conflict handling mechanism of $\mathcal{DL}^{\mathrm{N}}$ is useful also in other advocated applications of nonmonotonic reasoning. One of them is the mitigation of risks due to human errors. For example, contraindications may be associated to new drugs by default, unless explicitly stated otherwise, to prevent the risk of forgetting some of them. However, if a drug belongs to two or more drug families, then multiple inheritance may cause unresolvable conflicts, that prevent the correct contraindications from being inherited. If such conflicts were silently removed, then the missing contraindications might remain unnoticed, while in $\mathcal{DL}^{\mathrm{N}}$ the potential loss of crucial information would be signaled through a concept's inconsistency, by analogy with the above example. This application of nonmonotonic reasoning has been suggested in [36] and its encoding in $\mathcal{DL}^{\mathrm{N}}$ has been discussed in [5, Appendix C].

We proceed with another application example from the semantic policy domain, that shows $\mathcal{DL}^{\mathrm{N}}$'s behavior on multiple exception levels.

Example 2 We are going to axiomatize the following natural language policy: *"In general, users cannot access confidential files; Staff can read confidential files; Blacklisted users are not granted any access. This directive cannot be overridden."* Note that each of the above directives contradicts (and is supposed to override) its predecessor in some particular case. Authorizations can be reified as objects with attributes *subject* (the access requestor), *target* (the file to be accessed), and *privilege* (such as *read* and *write*). Then the above policy can be encoded as follows:

$$\mathtt{Staff} \sqsubseteq \mathtt{User} \tag{9}$$

$$\mathtt{Blklst} \sqsubseteq \mathtt{Staff} \tag{10}$$

$$\mathtt{UserReqst} \sqsubseteq_n \neg\exists\mathtt{privilege} \tag{11}$$

$$\mathtt{StaffReqst} \sqsubseteq_n \exists\mathtt{privilege.Read} \tag{12}$$

$$\mathtt{BlkReq} \sqsubseteq \neg\exists\mathtt{privilege} \tag{13}$$

where $\mathtt{BlkReq} = \exists\mathtt{subj.Blklst}$, $\mathtt{StaffReqst} = \exists\mathtt{subj.Staff}$, and $\mathtt{UserReqst} = \exists\mathtt{subj.User}$. By (9), both of the specifity relations yield $(12) \prec (11)$, that is, (12) has higher priority than (11) under priority (1) as well as (2). Let $\Sigma = \{\mathrm{N}\mathtt{StaffReqst}\}$; (12) overrides (11) in $\mathrm{N}\mathtt{StaffReqst}$ so $\mathcal{KB}^{\Sigma}$ consists of: (9), (10), (13), plus

$$\mathrm{N}\mathtt{StaffReqst} \sqsubseteq \mathtt{StaffReqst}$$

$$\mathrm{N}\mathtt{StaffReqst} \sqcap \mathtt{StaffReqst} \sqsubseteq \exists\mathtt{privilege.Read}\,.$$

Consequently, $\mathcal{KB} \models\!\!\approx \mathrm{N}\mathtt{StaffReqst} \sqsubseteq \exists\mathtt{privilege.Read}$. Similarly, it can be verified that:

1. Normally, access requests involving confidential files are rejected, if they come from generic users. Formally, $\mathcal{KB} \models\!\!\approx \mathrm{N}\mathtt{UserReqst} \sqsubseteq \neg\exists\mathtt{privilege}$;

2. Blacklisted users cannot do anything by (13), so, in particular: $\mathcal{KB} \models\!\!\approx \mathrm{N}\mathtt{BlkReq} \sqsubseteq \neg\exists\mathtt{privilege}$. ∎

Some application examples from the biomedical domain can be found in [5] (see Examples 3, 4, 10, 12, and the drug contraindication example in Appendix C). Like the above examples, they are all correctly solved by $\mathcal{DL}^{\mathrm{N}}$ with both priority notions. The examples inspired by applications hardly exhibit the complicated networks of dependencies between conflicting defaults that occur in artificial examples. Nonetheless, we briefly discuss such examples, too, as a means of comparing $\mathcal{DL}^{\mathrm{N}}$ with other logics such as [38, 22, 8].

In several cases, e.g. examples B.4 and B.5 in [38], $\mathcal{DL}^{\mathrm{N}}$ agrees with [38, 22, 8] under both priority relations. Since these examples have similar features, here we illustrate only B.4, as a representative case.

Example 3 (Juvenile offender) Let $\mathcal{KB}$ consist of axioms (14)–(18) where `J`, `G`, `M`, `P` abbreviate `JuvenileOffender`, `GuiltyOfCrime`, `IsMinor` and `ToBePunished`, respectively.

$$\mathtt{J} \sqsubseteq \mathtt{G} \quad (14)$$
$$\mathtt{J} \sqsubseteq \mathtt{M} \quad (15)$$
$$\mathtt{M} \sqcap \mathtt{G} \sqsubseteq_n \neg\mathtt{P} \quad (16)$$
$$\mathtt{M} \sqsubseteq_n \neg\mathtt{P} \quad (17)$$
$$\mathtt{G} \sqsubseteq_n \mathtt{P} \quad (18)$$

$$\mathtt{J} \sqsubseteq \mathtt{G} \quad (19)$$
$$\mathtt{J} \sqsubseteq \mathtt{M} \quad (20)$$
$$\mathrm{N}\mathtt{J} \sqsubseteq \mathtt{J} \quad (21)$$
$$\mathrm{N}\mathtt{J} \sqcap \mathtt{M} \sqcap \mathtt{G} \sqsubseteq \neg\mathtt{P} \quad (22)$$
$$\mathrm{N}\mathtt{J} \sqcap \mathtt{M} \sqsubseteq \neg\mathtt{P} \quad (23)$$

On one hand, criminals have to be punished and, on the other hand, minors cannot be punished. So, what about juvenile offenders? The defeasible inclusion (16) breaks the tie in favor of their being underage, hence not punishable. By setting $\Sigma = \{\mathrm{N}\mathtt{J}\}$, priorities (1) and (2) both return axioms (19)–(23) as $\mathcal{KB}^{\Sigma}$. Then, clearly, $\mathcal{KB}^{\Sigma} \models \mathrm{N}\mathtt{J} \sqsubseteq \neg\mathtt{P}$ which is $\mathcal{DL}^{\mathrm{N}}$'s analogue of the inferences of [38, 22, 8]. ∎

In other cases (e.g. example B.1 in [38]) $\mathcal{DL}^{\mathrm{N}}$ finds the same conflicts as [38, 22, 8]. However, $\mathcal{DL}^{\mathrm{N}}$'s semantics signals these conflicts to the knowledge engineer whereas in [38, 22, 8] they are silently neutralized.

Example 4 (Double Diamond) Let $\mathcal{KB}$ be the following set of axioms:

$$\mathtt{A} \sqsubseteq_n \mathtt{T} \quad (24)$$
$$\mathtt{A} \sqsubseteq_n \mathtt{P} \quad (25)$$
$$\mathtt{T} \sqsubseteq_n \mathtt{S} \quad (26)$$
$$\mathtt{P} \sqsubseteq_n \neg\mathtt{S} \quad (27)$$

$$\mathtt{S} \sqsubseteq_n \mathtt{R} \quad (28)$$
$$\mathtt{P} \sqsubseteq_n \mathtt{Q} \quad (29)$$
$$\mathtt{Q} \sqsubseteq_n \neg\mathtt{R} \quad (30)$$

DIs (26) and (27) have incomparable priority under (1) and (2). Consequently, it is easy to see that $\mathrm{N}\mathtt{A} \sqsubseteq \mathtt{S}$ and $\mathrm{N}\mathtt{A} \sqsubseteq \neg\mathtt{S}$ are both implied by $\mathcal{KB}^{\Sigma}$ and hence the knowledge engineer is warned that $\mathrm{N}\mathtt{A}$ is inconsistent. The same conflict is silently neutralized in [22, 38, 8] (`A`'s instances are subsumed by neither `S` nor $\neg$`S` and no inconsistency arises). Similarly for the incomparable DIs (28) and (30) and the related conflict. ∎

The third category of examples (e.g. B.2 and B.3 in [38]) presents a more variegated behavior. In particular, using priority (2), both [22] and $\mathcal{DL}^{\mathrm{N}}$ solve all conflicts and infer the same consequences; [38] solves only some conflicts; [8] is not able to solve any conflict and yet it does not raise any inconsistency warning; $\mathcal{DL}^{\mathrm{N}}$ with priority (1) cannot solve the conflicts but raises an inconsistency warning. Here, for the sake of simplicity, we discuss in detail a shorter example which has all relevant ingredients.

Example 5 Let $\mathcal{KB}$ be the following defeasible knowledge base:

$$\mathtt{A} \sqsubseteq_n \mathtt{B} \quad (31) \qquad \mathtt{A} \sqsubseteq_n \mathtt{C} \quad (32) \qquad \mathtt{B} \sqsubseteq_n \neg\mathtt{C} \quad (33)$$

According to priority (1) all DIs are incomparable. Therefore, $\mathcal{DL}^{\mathrm{N}}$ warns (by inferring $\mathrm{N}\mathtt{A} \sqsubseteq \bot$) that the conflict between $\mathrm{N}\mathtt{A} \sqsubseteq \mathtt{C}$ and $\mathrm{N}\mathtt{A} \sqsubseteq \neg\mathtt{C}$ cannot be solved. Note that [8] adopts priority (1), too, however according to circumscription, any interpretation where $\mathtt{A}$'s instances are either in $\neg\mathtt{C} \sqcap \mathtt{B}$ or in $\mathtt{C}$ is a model, so $\mathtt{A}$ is satisfiable (the conflict is silently neutralized). Under priority (2), instead, axiom (31) gives (31) itself and (32) higher priority than (33). Consequently, $\mathrm{N}\mathtt{A} \sqsubseteq \mathtt{C}$ prevails over $\mathrm{N}\mathtt{A} \sqsubseteq \neg\mathtt{C}$. In this case, $\mathcal{DL}^{\mathrm{N}}$ and rational closure infer the same consequences. ∎

There is a reason why the inferences of $\mathcal{DL}^{\mathrm{N}}$ are so similar to those of [22] in the above examples. We will show in the following sections that $\mathcal{DL}^{\mathrm{N}}$ and [22] satisfy similar versions of a set of standard postulates (called *KLM postulates*) in all the examples where the characteristic features of $\mathcal{DL}^{\mathrm{N}}$ do not come into play. This is why the above examples show the different behavior of $\mathcal{DL}^{\mathrm{N}}$ only in the presence of unresolved conflicts. Further differences will emerge in Section 5.2 from the restriction of role values to normal individuals – a feature that is not considered in the above examples and that is not supported by [22].

2.3 The KLM postulates

In their seminal papers, Makinson [34] and Kraus, Lehmann, and Magidor [29, 32, 30] study the consequence relations of nonmonotonic logics from an axiomatic perspective. In order to develop a general framework, compatible with the variety of nonmonotonic languages and semantics introduced in the literature, they abstract away the details of nonmonotonic expressions and focus on the classical information that such expressions deduce from the available classical description of the world.

For example, consider Reiter's default theories [37]: they are pairs $\langle D, W\rangle$ where D is a set of default rules and W is a set of classical sentences. Every set of defaults D is associated to the consequence relation $Cn_D(\cdot)$ that maps each possible W on the set of classical sentences that can be derived from W using the default rules in D [34]. In later works, an equivalent representation in terms of sequents is adopted. For instance, $Cn_D(\cdot)$ can be equivalently represented as the set of all pairs of sentences (ω, ϕ) such that ω is the conjunction of W's sentences, and $\phi \in Cn_D(W)$. Such sequents are usually denoted by $\omega \mid\!\sim \phi$. Note that $\mid\!\sim$ embodies the background knowledge represented by D, hiding both its semantic contents and its syntactic representation.

Kraus, Lehmann, and Magidor argued that in order to reason about what normally holds in the world, it is desirable to make nonmonotonic consequence relations closed under certain properties, called *KLM postulates*, from the initials of their authors. The verbatim instantiation of the original postulates, using $\mathcal{DL}^{\mathrm{N}}$'s terminology, is illustrated in Table 1. Here, the background knowledge hidden in $\mid\!\sim$ is made explicit by the term $\mathcal{KB}$, and each entailment of the form $\mathcal{KB} \cup \{\alpha\} \models\!\!\!\approx \beta$ corresponds to the sequent $\alpha \mid\!\sim \beta$, by analogy with the above example about default logic. Through this correspondence, one gets precisely the original postulates. A consequence relation that

Table 1: The KLM postulates in $\mathcal{DL}^{\mathrm{N}}$

Name	Rule schema	Sound in $\mathcal{DL}^{\mathrm{N}}$
REF	$\dfrac{\alpha \in \mathcal{KB}}{\mathcal{KB} \mathrel{\mid\!\approx} \alpha}$	✓
CT	$\dfrac{\mathcal{KB} \mathrel{\mid\!\approx} \alpha \quad \mathcal{KB} \cup \{\alpha\} \mathrel{\mid\!\approx} \gamma}{\mathcal{KB} \mathrel{\mid\!\approx} \gamma}$	✓
CM	$\dfrac{\mathcal{KB} \mathrel{\mid\!\approx} \alpha \quad \mathcal{KB} \mathrel{\mid\!\approx} \gamma}{\mathcal{KB} \cup \{\alpha\} \mathrel{\mid\!\approx} \gamma}$	✓
LLE	$\dfrac{\mathcal{KB} \cup \{\alpha\} \mathrel{\mid\!\approx} \gamma \quad \models \alpha \equiv \beta}{\mathcal{KB} \cup \{\beta\} \mathrel{\mid\!\approx} \gamma}$	✓
RW	$\dfrac{\mathcal{KB} \mathrel{\mid\!\approx} \alpha \quad \alpha \models \gamma}{\mathcal{KB} \mathrel{\mid\!\approx} \gamma}$	✓
OR	$\dfrac{\mathcal{KB} \cup \{\alpha\} \mathrel{\mid\!\approx} \gamma \quad \mathcal{KB} \cup \{\beta\} \mathrel{\mid\!\approx} \gamma}{\mathcal{KB} \cup \{\alpha \vee \beta\} \mathrel{\mid\!\approx} \gamma}$	
RM	$\dfrac{\mathcal{KB} \mathrel{\mid\!\approx} \gamma \quad \mathcal{KB} \not\mathrel{\mid\!\approx} \neg\alpha}{\mathcal{KB} \cup \{\alpha\} \mathrel{\mid\!\approx} \gamma}$	

$\mathcal{KB}$ is a canonical $\mathcal{DL}^{\mathrm{N}}$ knowledge base;
α and β range over $\mathcal{DL}$ assertions and (strong) concept and role inclusions;
γ ranges over $\mathcal{DL}^{\mathrm{N}}$ assertions and $\mathcal{DL}^{\mathrm{N}}$ concept/role inclusions;
nonstandard DL axioms $\alpha \vee \beta$, $\neg\beta$ can be simulated, e.g. with the universal role;
$\mathrel{\mid\!\approx}$ denotes the nonmonotonic consequence relation of $\mathcal{DL}^{\mathrm{N}}$ and $\models$ denotes classical inference.

satisfies the KLM postulates is called *rational*. It is called *preferential* if it satisfies all rules but RM, and *cumulative* if it satisfies all rules but RM and OR.[8]

In their abstract analysis of nonmonotonic consequence relations, Kraus, Lehmann and Magidor found it helpful to identify nonmonotonic theories with the set of sequents $\omega \mathrel{\mid\!\sim} \phi$ that constitute their consequence relations, so as to make their arguments independent from any concrete syntax for defeasible information. Then, some authors found it natural to turn these sequents (that originally were metalevel expressions that describe the effects of a theory) into object-level expressions that *are* the theory itself, whose semantics can be specified – say – by the KLM postulates and their model-theoretic accounts, or by the rational closure construction. Some examples of this approach in DLs are [20, 21, 22, 19, 16]. This *internalization* operation transformed the KLM postulates from a general analysis tool – applicable to any logic – into a specific logic, just like default logic, autoepistemic logic or circumscription.

According to the above discussion, we shall use the metalevel version reported in Table 1 as the analysis tool to study $\mathrel{\mid\!\approx}$, i.e. the consequence relation of $\mathcal{DL}^{\mathrm{N}}$, and compare it with those of Default, Autoepistemic, and Circumscribed DLs. This metalevel analysis will be refined and extended by using two internalized versions of the

[8]Kraus, Lehmann and Magidor study also further postulates, including one called LOOP that cannot be derived from the postulates of a cumulative consequence relation. The study of these postulates in the context of $\mathcal{DL}^{\mathrm{N}}$ lies beyond the scope of this paper and is left for further work.

KLM postulates in order to clarify the meaning of $\mathcal{DL}^{\mathrm{N}}$ expressions such as $C \sqsubseteq_n D$ and $\mathrm{N}C \sqsubseteq D$, that syntactically resemble the internalized conditionals $\omega \mid\!\sim \phi$ but actually have different semantics. The internalized postulates help also in comparing $\mathcal{DL}^{\mathrm{N}}$ with Lehmann's default logic and with the description logics illustrated in the following section.

2.4 Two families of DLs satisfying internalized KLM postulates

Along the following sections, we shall refer to two families of nonmonotonic DLs that provide a natural reference point for comparison and illustrate different ways of supporting internalized KLM postulates. These logics have been selected precisely because they have been designed to satisfy such postulates (or most of them), unlike circumscribed and autoepistemic DLs, just to name a few.

The first family is obtained by applying the rational closure construction or variants thereof to $\mathcal{ALC}$, using sequents $C \mid\!\sim D$ as knowledge base axioms that informally mean: "*normally, the instances of C satisfy D*" [20, 21, 22, 19, 16].[9] The priority relation over these axioms is given by the rank computed by Algorithm 1. The symbol $\mid\!\sim$ is part of the object-level language, and can be used both in knowledge bases and in queries. Unlike $\mathcal{DL}^{\mathrm{N}}$, there are no constructs for denoting the normal instances of a concept. We shall refer to these logics as *DLs based on Rational Closure*. In these logics role values cannot be restricted to normal individuals; equivalently, defeasible assertions do not apply to role values. Recently, a dual logic where roles range on normal individuals but inclusions are classical has been studied in [14]. Most DLs based on rational closure satisfy the following internalized version of the KLM postulates:

(REF$_{RC}$) $\quad C \mid\!\sim C$

(LLE$_{RC}$) $\quad \dfrac{C \mid\!\sim E \quad \models C \equiv D}{D \mid\!\sim E}$

(RW$_{RC}$) $\quad \dfrac{C \mid\!\sim D \quad \models D \sqsubseteq E}{C \mid\!\sim E}$

(RM$_{RC}$) $\quad \dfrac{C \mid\!\sim E \quad C \not\mid\!\sim \neg D}{C \sqcap D \mid\!\sim E}$

(CT$_{RC}$) $\quad \dfrac{C \mid\!\sim D \quad C \sqcap D \mid\!\sim E}{C \mid\!\sim E}$

(CM$_{RC}$) $\quad \dfrac{C \mid\!\sim D \quad C \mid\!\sim E}{C \sqcap D \mid\!\sim E}$

(OR$_{RC}$) $\quad \dfrac{C \mid\!\sim E \quad D \mid\!\sim E}{C \sqcup D \mid\!\sim E}$

The only exception is the relevant closure introduced in [19], that in order to prevent inheritance blocking waives OR$_{RC}$, CM$_{RC}$, and RM$_{RC}$.

The second family, that we call *typicality DLs*, feature concepts $T(C)$ that denote the typical instances of C; they are analogues of the normality concepts $\mathrm{N}C$. The underlying monotonic logic is essentially a preferential modal logic whose semantics is based on a normality relation over individuals, that is illustrated below. Typicality DLs have no ad hoc construct for defeasible inclusions; the analogue of $C \sqsubseteq_n D$ and $C \mid\!\sim D$ is $T(C) \sqsubseteq D$. So, strictly speaking, in typicality logics the inclusions themselves are not defeasible; a statement like $T(C) \sqsubseteq D$ cannot be overridden. Nonmonotonicity is obtained by minimizing atypical individuals, by a metalevel construction similar to circumscription [24, 25]. Interestingly, the expressions $T(C) \sqsubseteq D$ are

[9] In some of these papers, the symbol $\mid\!\sim$ is replaced by $\sqsubseteq\!\!\!\sim$. Note that differently from the classical KLM framework, here C and D are concept expressions.

also syntactically similar to $\mathcal{DL}^{\mathrm{N}}$'s inclusions $\mathrm{N}C \sqsubseteq D$; their differences will be illustrated in Section 5. The logic $\mathcal{ALC} + \mathbf{T}_{min}$ illustrated in [24] is preferential, while those in [25] are rational, relative to the following internalized KLM postulates:

$$(\mathrm{REF}_T) \quad T(C) \sqsubseteq C \qquad\qquad (\mathrm{CT}_T) \quad \frac{T(C) \sqsubseteq D \quad T(C \sqcap D) \sqsubseteq E}{T(C) \sqsubseteq E}$$

$$(\mathrm{LLE}_T) \quad \frac{T(C) \sqsubseteq E \quad \models C \equiv D}{T(D) \sqsubseteq E} \qquad\qquad (\mathrm{CM}_T) \quad \frac{T(C) \sqsubseteq D \quad T(C) \sqsubseteq E}{T(C \sqcap D) \sqsubseteq E}$$

$$(\mathrm{RW}_T) \quad \frac{T(C) \sqsubseteq D \quad \models D \sqsubseteq E}{T(C) \sqsubseteq E} \qquad\qquad (\mathrm{OR}_T) \quad \frac{T(C) \sqsubseteq E \quad T(D) \sqsubseteq E}{T(C \sqcup D) \sqsubseteq E}$$

$$(\mathrm{RM}_T) \quad \frac{T(C) \sqsubseteq E \quad T(C) \not\sqsubseteq \neg D}{T(C \sqcap D) \sqsubseteq E}$$

Many of the logical properties of typicality DLs are forced by their monotonic fragments, such as $\mathcal{ALC} + \mathbf{T}$ and $\mathcal{ALC} + \mathbf{T_R}$. In each model $\mathcal{I}$, the degree of typicality is formalized with a partial order $\leq^{\mathcal{I}}$ over individuals. The typical instances of a concept C, denoted by $T(C)$, are the $\leq^{\mathcal{I}}$-minimal instances of C, that is, its most typical members. Different typicality logics are obtained by tuning the properties of the relations $\leq^{\mathcal{I}}$. They are always assumed to satisfy a condition called *smoothness* that ensures that each nonempty concept has at least one $\leq^{\mathcal{I}}$-minimal instance; in other words, if C is consistent, then $T(C)$ is consistent, too. Note that $\leq^{\mathcal{I}}$ does not depend on C, that is, a single, universal normality criterion is assumed. Informally speaking, this means that if John is more typical than Mary as a parent, then he must also be more typical than Mary as a driver, as a worker, as a tax payer, and so on. Some technical and practical consequences of this strong assumption are discussed in the following sections, mostly in Section 5.2.

3 $\mathcal{DL}^{\mathrm{N}}$ between cumulativity and rationality

In this section we prove that the entailment relation of $\mathcal{DL}^{\mathrm{N}}$ (that is $\models\!\!\!\approx$) is cumulative, and that if normal instances are assumed to exist, then it becomes rational. In some logics this assumption is hardwired in the underlying monotonic logic, while in $\mathcal{DL}^{\mathrm{N}}$ it must be explicitly stated through suitable axioms. This is related to $\mathcal{DL}^{\mathrm{N}}$'s strategy for avoiding undesired closed-world effects; in particular, $\mathcal{DL}^{\mathrm{N}}$ does not assume that any individual is normal by default, see [5] for an articulated discussion of this approach.

3.1 $\mathcal{DL}^{\mathrm{N}}$ is cumulative

Obviously, REF and RW always hold, because $\mathcal{DL}^{\mathrm{N}}$ is closed under classical inference. Moreover:

Theorem 1 *Rules CT, CM, and LLE are sound.*

Since $\models\!\!\!\approx$ satisfies REF, CT, CM, LLE, and RW, it is a *cumulative consequence relation*. The other rules, OR and RM, do not universally hold. In the next subsections we analyze why and show that they actually hold under a simple additional condition.

3.2 The RM rule

This postulate can be analyzed through an extremely simple example: given the knowledge base

$$\mathcal{KB} = \{A \sqsubseteq_n B\} ,$$

there is little doubt that all the standard instances of A should belong to B, that is,

$$\mathcal{KB} \mathrel{|\!\!\approx} \mathrm{N}A \sqsubseteq B . \tag{34}$$

There should also be little doubt that the strong axiom $A \sqsubseteq \neg B$ overrides the DI in $\mathcal{KB}$, so that

$$\mathcal{KB} \cup \{A \sqsubseteq \neg B\} \not\mathrel{|\!\!\approx} \mathrm{N}A \sqsubseteq B . \tag{35}$$

Given (34) and (35), RM can be satisfied only if

$$\mathcal{KB} \mathrel{|\!\!\approx} \neg(A \sqsubseteq \neg B) . \tag{36}$$

Now the question is: why should (36) hold? $\mathcal{KB}$ only says that the *normal* instances of A are in B. Then, concluding that A is not included in $\neg B$ is like assuming that at least one of such instances exists.

The assumption that each consistent concept has at least one normal instance is hardwired in the semantics of typicality logics, through smoothness.[10] On the contrary, $\mathcal{DL}^{\mathrm{N}}$ does not make this assumption, and for this reason RM does not always hold. This is witnessed by the above example, since (36) does not hold in $\mathcal{DL}^{\mathrm{N}}$. If normal instances were assumed to exist, then $\mathcal{DL}^{\mathrm{N}}$ would satisfy RM, as well:

Theorem 2 *Consider RM and the axiom γ occurring in it. Suppose that $\mathcal{KB}$ contains an axiom $\neg(\mathrm{N}C \sqsubseteq \bot)$ (or equivalent formulations, cf. Sec. 2.1) for each $\mathrm{N}C$ occurring in $\mathcal{KB} \cup \{\gamma\}$ such that C is satisfiable w.r.t. $\mathcal{KB}$. Then RM is sound.*

Remark 1 Note that the knowledge base, in the above theorem, is not canonical; it is a *general* knowledge base where normality concepts occur in the strong part. In general, such unrestricted use of normality concepts may lead to undesirable inferences, by overriding the intended meaning carried by defeasible inclusions – which is why we focussed on the canonical approach, cf. the note on usage in Section 2.2. However, if non-canonical axioms are carefully chosen, then non-canonical knowledge bases may become a useful tool for specializing $\mathcal{DL}^{\mathrm{N}}$ and tuning the set of valid axioms to application needs (more examples will be given in the following sections). This approach does not have any major side-effects on the properties of $\mathcal{DL}^{\mathrm{N}}$ since they all rely on the correctness of the translation $\mathcal{KB}^{\Sigma}$, which has been proved in [5] for *all* general knowledge bases, not just canonical ones. The selective use of general knowledge bases as a means to achieve more flexibility is an interesting topic for further research. ▮

The analysis of RM's effects on inconsistent prototypes is illuminating, too. Consider a standard case of unresolved conflict arising from multiple inheritance:

[10] Smoothness was first used in [29], where it was applied to possible worlds, though. Interestingly, a similar assumption is made in [31] to satisfy the postulates for $\exists$.

$$A_1 \sqsubseteq_n B \quad (37) \qquad C \sqsubseteq A_1 \sqcap A_2 \quad (39)$$
$$A_2 \sqsubseteq_n \neg B \quad (38)$$

In the above $\mathcal{KB}$, the prototype of C is inconsistent, i.e.

$$\mathcal{KB} \approx\!\!\!\models \mathrm{N}C \sqsubseteq \bot\,, \quad (40)$$

due to the conflicting DIs (37) and (38). The natural way of resolving such conflicts is specifying which alternative should hold for C, thereby overriding the other one. For instance, in our example, one might specify that $C \sqsubseteq B$ and expect that

$$\mathcal{KB} \cup \{C \sqsubseteq B\} \not\approx\!\!\!\models \mathrm{N}C \sqsubseteq \bot \quad (41)$$

because $C \sqsubseteq B$ clearly overrides (38). Now consider an instance of RM where $\alpha = (C \sqsubseteq B)$ and $\gamma = (\mathrm{N}C \sqsubseteq \bot)$. Given (40) and (41), that are the first premise and the negation of the conclusion, in order to satisfy RM, it is necessary that the second premise does not hold, that is,

$$\mathcal{KB} \approx\!\!\!\models \neg(C \sqsubseteq B)\,, \quad (42)$$

and again the question is: why should (42) hold? Differently from the previous example, here $\mathcal{KB}$ entails that $\mathrm{N}C$ is empty, so the instances of C that belong to $\neg B$ (that must exist if (42) holds) are *not* typical instances of C. It follows that enforcing RM in a logic like $\mathcal{DL}^{\mathrm{N}}$, that highlights unresolved conflicts through inconsistent prototypes, requires extending the effects of DIs from normal instances only to *all* individuals. We leave it as an open question whether this is appropriate, and how it should be done.

3.3 The OR rule

Similarly to RM, the OR rule depends on the assumption that normal individuals exist. To see this, we first illustrate a counterexample to OR. Let $\mathcal{KB}$ be the following set of axioms:

$$A \sqsubseteq_n \{c\} \quad (43) \qquad A \sqsubseteq B \quad (46)$$
$$A \sqsubseteq_n \{d\} \quad (44) \qquad B \sqsubseteq \exists R \quad (47)$$
$$c \neq d \quad (45) \qquad B \sqsubseteq_n \forall R.(\mathrm{N}A \sqcap C) \quad (48)$$

and let

$$\alpha = A \sqsubseteq \{c\}\,,$$
$$\beta = A \sqsubseteq \{d\}\,,$$

Note that $\alpha \vee \beta$ is equivalent to $A \sqsubseteq \{c, d\}$, so it can be expressed in standard description logics. Here α cleary overrides (44) and β overrides (43) , so the prototype of A

is consistent in both $\mathcal{KB} \cup \{\alpha\}$ and $\mathcal{KB} \cup \{\beta\}$:

$$\begin{aligned}\mathcal{KB} \cup \{\alpha\} &\not\models\!\!\approx \quad \mathrm{N}A \sqsubseteq \bot\,,\\ \mathcal{KB} \cup \{\beta\} &\not\models\!\!\approx \quad \mathrm{N}A \sqsubseteq \bot\,.\end{aligned}$$

Since $\mathrm{N}A$ is nonempty, (48) is not overridden and we have:

$$\begin{aligned}\mathcal{KB} \cup \{\alpha\} &\models\!\!\approx \quad \mathrm{N}B \sqsubseteq \forall R.C\,,\\ \mathcal{KB} \cup \{\beta\} &\models\!\!\approx \quad \mathrm{N}B \sqsubseteq \forall R.C\,.\end{aligned}$$

From these entailments, the OR rule derives $\mathcal{KB} \cup \{\alpha \vee \beta\} \models\!\!\approx \mathrm{N}B \sqsubseteq \forall R.C$. However, this does not hold in $\mathcal{DL}^{\mathrm{N}}$. Axiom $\alpha \vee \beta$ is too weak to override any of (43) and (44), so their conflict remains unresolved and we have

$$\mathcal{KB} \cup \{\alpha \vee \beta\} \models\!\!\approx \mathrm{N}A \sqsubseteq \bot\,.$$

As a consequence, (48) is overridden (as it would be inconsistent with (47)) and hence

$$\mathcal{KB} \cup \{\alpha \vee \beta\} \not\models\!\!\approx \mathrm{N}B \sqsubseteq \forall R.C\,.$$

Note that this counterexample exploits an empty normality concept. This is not by chance: similarly to RM, we can prove that $\mathcal{DL}^{\mathrm{N}}$ satisfies OR when normal individuals are guaranteed to exist.

Theorem 3 *Consider OR and the axiom γ occurring in it. Suppose that $\mathcal{KB}$ contains an axiom $\neg(\mathrm{N}C \sqsubseteq \bot)$ (or equivalent formulations) for each $\mathrm{N}C$ occurring in $\mathcal{KB} \cup \{\gamma\}$ such that C is satisfiable w.r.t. $\mathcal{KB}$. Then OR is sound.*

Table 1 will be used to compare $\mathcal{DL}^{\mathrm{N}}$ with other nonmonotonic DLs in sections 7 and 8.

4 The logical properties of DIs

From a crudely syntactic point of view, a defeasible inclusion $C \sqsubseteq_n D$ is similar to its counterpart $C \mid\!\sim D$ adopted by the DLs based on rational closure. So it may be tempting to assume that DIs should satisfy the same logical properties, including the KLM postulates. The corresponding internalized version is reported in Table 2.[11]

The semantics of DIs, however, is different: while the expressions $C \mid\!\sim D$ are rational sequents, that are required to belong to the knowledge base's consequences, DIs express only *default assertions* that may be overridden and have no effect. In this respect, DIs are more similar to Reiter's default rules [37] and Lehmann's version thereof [30]; in particular, in Lehmann's terms, our DIs correspond to the default

[11]The intended meaning of these inference rules is that if $\mathcal{KB}$ entails (with $\models\!\!\approx$) the DIs in the premises, then it should entail the conclusions. Non-defeasible premises are interpreted in the obvious way.

Table 2: Analogues of the KLM postulates for DIs [5]

Name	Rule schema	Sound in $\mathcal{DL}^{\mathrm{N}}$
REF_n	$\dfrac{}{C \sqsubseteq_n C}$	✓
CT_n	$\dfrac{C \sqsubseteq_n D \quad C \sqcap D \sqsubseteq_n E}{C \sqsubseteq_n E}$	
CM_n	$\dfrac{C \sqsubseteq_n D \quad C \sqsubseteq_n E}{C \sqcap D \sqsubseteq_n E}$	
LLE_n	$\dfrac{C \sqsubseteq_n E \quad \mathcal{S} \models C \equiv D}{D \sqsubseteq_n E}$	partly
RW_n	$\dfrac{C \sqsubseteq_n D \quad \mathcal{S} \models D \sqsubseteq E}{C \sqsubseteq_n E}$	
OR_n	$\dfrac{C \sqsubseteq_n E \quad D \sqsubseteq_n E}{C \sqcup D \sqsubseteq_n E}$	partly
RM_n	$\dfrac{C \sqsubseteq_n E \quad C \not\sqsubseteq_n \neg D}{C \sqcap D \sqsubseteq_n E}$	

$\mathcal{S}$ is the strong part of the knowledge base expressions $(\omega : \phi)$, not to the sequents $\omega \mid\sim \phi$ that represent the consequences of default expressions. Default rules have the following characteristic behavior: either they are overridden (and have no effect whatsoever, like tautologies) or they are applied; in the latter case, defaults have a pretty classical behavior (e.g. Lehmann's are equivalent to material implications). Indeed, if this aspect were not properly taken into account by distinguishing so-called "meaningful" (i.e. applied) defaults from overridden ones, then even Lehmann's account of default reasoning – which is built around the KLM postulates – would not be strictly rational; see the discussion in [30, Sec. 5 and 6].

The analogue of these properties in $\mathcal{DL}^{\mathrm{N}}$ is cast into the semantics of DIs: δ is satisfied if for all concepts $\mathrm{N}C$, either δ is overridden in $\mathrm{N}C$ or the corresponding classical inclusion $\delta^{\mathrm{N}C}$ holds. In particular, if δ is overridden in every $\mathrm{N}C$, then it is satisfied even if it has no effects.

Said so, it comes to no surprise that almost none of the rules in Table 2 is unconditionally sound. For example, the conclusion of RW_n is justified by its first premise; if it were overridden, then the conclusion would not be supported. Moreover, the conclusion is weaker than the premise, so it is not necessarily overridden when the premise is. Rule CT_n may even produce undesirable effects if it were forced in general (cf. Theorem 8 and the following discussion in [5]).

However, as it should be expected, in line with the characteristic features of default reasoning, *when the premises are not overridden, then all of the rules in Table 2 are sound* [5, Theorem 9]. In particular, REF_n is always sound.

Actually, under pretty mild assumptions, some rules are sound also when their defeasible premises are overridden. Rule LLE_n, for example, can be violated only by using bizarre priority relations, that depend on syntactic details and give logically equivalent DIs different priority. If the priority relation is not sensitive to syntactic details, then LLE_n holds [5, Theorem 10]. In [5, Theorem 8], it has also been proved that OR_n holds unconditionally under specificity, because this priority relation guarantees that the conclusion is overridden whenever both premises are overridden, thereby ensuring the rule's soundness.

Given that DIs are like defaults, it is interesting to compare $\mathcal{DL}^{\mathrm{N}}$'s consequence relation with those of the other default logics; this will be done in detail in sections 7 and 8. We informally anticipate that $\mathcal{DL}^{\mathrm{N}}$ (which is cumulative) lies somewhere in between Default DLs (which are not cumulative) and Lehmann's propositional default logic, which is rational [30]. With the additional assumption that consistent concepts have standard instances, the consequence relation of $\mathcal{DL}^{\mathrm{N}}$ is rational, like that of Lehmann's default logic.

5 The logical properties of N

5.1 N and boolean operators

Normality concepts satisfy the following natural axiom schema:

$$\mathrm{N}C \sqsubseteq C \tag{49}$$

Table 3: Candidate axioms relating N with boolean operators

Name	Axiom schema	Sound in $\mathcal{DL}^{\mathrm{N}}$
neg 1	$\mathrm{N}\neg C \sqsubseteq \neg\mathrm{N}C$	✓
neg 2	$\mathrm{N}\neg C \sqsupseteq \neg\mathrm{N}C$	
and 1	$\mathrm{N}C \sqcap \mathrm{N}D \sqsubseteq \mathrm{N}(C \sqcap D)$	
and 2	$\mathrm{N}C \sqcap \mathrm{N}D \sqsupseteq \mathrm{N}(C \sqcap D)$	
or 1	$\mathrm{N}C \sqcup \mathrm{N}D \sqsubseteq \mathrm{N}(C \sqcup D)$	
or 2	$\mathrm{N}C \sqcup \mathrm{N}D \sqsupseteq \mathrm{N}(C \sqcup D)$	

C and D range over $\mathcal{DL}$ concepts

which is *strong*, that is, it cannot be overridden. More candidate axioms, relating normality concepts with the standard boolean operators, are listed in Table 3. However, not all of them make sense.

Axiom **neg 1** follows from schema (49) and its contrapositive:

$$\mathrm{N}\neg C \sqsubseteq \neg C \sqsubseteq \neg\mathrm{N}C\,.$$

It is easy to see that **neg 2** is undesirable, instead. The contrapositives of **neg 2** and (49) yield:

$$\mathrm{N}C \sqsupseteq \neg\mathrm{N}\neg C \sqsupseteq \neg\neg C \equiv C\,.$$

Together with (49) this would imply $\mathrm{N}C \equiv C$, that is, there could be no exceptional individuals, and there would be no difference between defeasible inclusions and strong axioms.

The other axioms for $\sqcap$ and $\sqcup$ violate the goal of representing exceptions, too. Many motivational examples in the literature contradict **and 2** and **or 1**. Consider, for instance, the semantic web policy where the authorizations for $\mathrm{N}\mathtt{Staff}$ and those for $\mathrm{N}\mathtt{Blacklisted}$ are mutually inconsistent (by default, everything is permitted to staff members while everything is strongly denied to blacklisted users). Then $\mathrm{N}\mathtt{Staff} \sqcap \mathrm{N}\mathtt{Blacklisted}$ is inconsistent, so **and 2** would make the concept $\mathrm{N}(\mathtt{Staff} \sqcap \mathtt{Blacklisted})$ inconsistent, too. On the contrary, the expected behavior is that the strong negative authorizations of blacklisted members override the default authorizations granted to staff members, so $\mathrm{N}(\mathtt{Staff} \sqcap \mathtt{Blacklisted})$ should be consistent. Next, consider the instance of **or 1** where C is $\mathtt{Staff}$ and D is $\mathtt{User}$ (recall that $\mathtt{Staff} \sqsubseteq \mathtt{User}$). Rule **or 1** forces $\mathrm{N}(\mathtt{Staff} \sqcup \mathtt{User})$ – whose instances should not have any access rights – to contain N`Staff`, whose instances *do* have read permission. Neither **and 2** nor **or 1** are valid in $\mathcal{DL}^{\mathrm{N}}$.

Axioms **and 1** and **or 2** are more controversial. They are valid in some logics, but we believe that a flexible nonmonotonic logic should not always satisfy them, for the reasons explained below.[12]

Example 6 (Drawbacks of and 1) Consider a country where the following situation holds. Most students (denoted by concept S) are not older than 25. Most employees (denoted by E) have a family income that exceeds € 1000 per month. Most working students ($S \sqcap E$) are older than 25, being slowed down by their job, and their family income is less than € 1000 per month. It is natural to assert that, by default, the instances of $\mathrm{N}S$ are not older than 25, that the family income of the instances of $\mathrm{N}E$ is at least € 1000, and that the instances of $\mathrm{N}(S \sqcap E)$ are older than 25 and their family income is less than € 1000. Note that the default properties of $\mathrm{N}(S \sqcap E)$, in this case, are incompatible with those of $\mathrm{N}S$ and $\mathrm{N}E$, so this natural approach at encoding the example is not compatible with **and 1**, that would force the instances of $\mathrm{N}(S \sqcap E)$ to contain the members of $\mathrm{N}S \sqcap \mathrm{N}E$, that are younger than 25 and have a family income grater than € 1000.[13] ∎

Example 7 (Drawbacks of or 2) A knowledge engineer shall encode the results of a marketing study in a knowledge base. Suppose that most people younger than 40 (denoted by concept Y) would drive a sport car, most people between 40 and 60 (denoted by M) would drive an agile city car, and most people over 60 (denoted by E) would drive a comfortable luxury car, while the rest of people would opt for a mid-sized car with low fuel consumption. To make calculations simpler, assume that in each case the majority amounts to 60% of the group, that people are equally distributed across the three age ranges, and that the four categories of cars are mutually disjoint. It seems natural to encode this knowledge by asserting that the instances of $\mathrm{N}Y$, $\mathrm{N}M$, and $\mathrm{N}E$ would choose by default a sport car, a city car or a luxury car, respectively. If we look at all people, though, (i.e. $Y \sqcup M \sqcup E$), we see that mid-sized cars get the largest share (40%, that is twice as large as any other choice, and *ceteris paribus* this ratio would grow further if there were more than 3 user categories). Therefore, the knowledge engineer may want to factorize common preferences and reduce the number of explicit car choice assertions in $\mathcal{KB}$ by selecting mid-sized cars with low fuel consumption as the default choice for the instances of $\mathrm{N}(Y \sqcup M \sqcup E)$. We believe that a flexible logic should allow this approach. On the contrary, this is not permitted by **or 2**: two applications of this axiom force the instances of $\mathrm{N}(Y \sqcup M \sqcup E)$ to belong to $\mathrm{N}Y \sqcup \mathrm{N}M \sqcup \mathrm{N}E$, and hence to conform to some of the typical choices of $\mathrm{N}Y$, $\mathrm{N}M$, and $\mathrm{N}E$ (which exclude mid-sized cars). Note that the majority of people indeed prefer a sport car or a city car or a luxury car, but this *disjunctive* information may be of little use in attributing automatically a *specific* default choice to the generic instances of $\mathrm{N}(Y \sqcup M \sqcup E)$, which explains the decision of our engineer. Choosing a mid-sized car by default, in

[12] Of course, some instances of **and 1** and **or 2** are not problematic, e.g. when $C \equiv D$.

[13] It is not hard to instantiate this scenario in such a way that $\mathrm{N}S \sqcap \mathrm{N}E$ is nonempty, so that the inclusion of this concept in $\mathrm{N}(S \sqcap E)$ manifests its undesirable effects.

this case, is an instance of the use of nonmonotonic reasoning as an opportunistic *storage convention*, aimed at reducing the size of the knowledge base, which is one of the possible uses of nonmonotonic reasoning envisioned by McCarthy [35]. ∎

In typicality DLs, **and 1** and **or 2** are strong axioms, enforced by the underlying monotonic logic $\mathcal{ALC}+\mathbf{T}$. It is easy to see (cf. Appendix B) that they are valid because the models of typicality logics adopt a single, concept-independent normality relation (e.g. if John is more typical than Mary as a student, it must also be more typical than Mary as an employee and as a working student). In $\mathcal{DL}^{\mathrm{N}}$, instead, each concept may be associated to an independent typicality criterion, which is the reason why **and 1** and **or 2** do not hold and the above two examples can be naturally encoded.

Remark 2 Note that **and 1** and **or 2** could be enforced in $\mathcal{DL}^{\mathrm{N}}$, if needed in some application, by including them in the knowledge base. A similar knowledge base is not canonical; it is a *general* knowledge base where normality concepts occur in the strong part. As we pointed out in Remark 1, this can be an interesting method for specializing $\mathcal{DL}^{\mathrm{N}}$ and tuning the set of its valid axioms when needed, so as to obtain maximum flexibility. ∎

5.2 Normality criteria and internalized KLM postulates

The inclusions $\mathrm{N}C \sqsubseteq D$ are analogues of the expressions $C \mid\!\sim D$ and $T(C) \sqsubseteq D$ when they are used as *queries*.

Recall from Section 2.4 that the models $\mathcal{I}$ of typicality logics are based on a universal, concept-independent notion of normality $\leq^{\mathcal{I}}$, and that $T(C)^{\mathcal{I}}$ is the set of $\leq^{\mathcal{I}}$-minimal members of $C^{\mathcal{I}}$, It follows easily from this definition that the monotonic fragment of typicality DLs satisfies the following inference rule:[14]

$$\frac{T(C) \sqsubseteq D}{T(C) \equiv T(C \sqcap D)} \ . \tag{50}$$

This rule clearly implies that CT_T and CM_T are sound. Since it may seem unrealistic to assume that the criterion for deciding what is a normal person is the same criterion that establishes – say – what is a normal hacker, in $\mathcal{DL}^{\mathrm{N}}$ each concept may have its own notion of what is more normal or standard. In technical terms, this implies that the counterpart of (50) does not hold, that is, $\mathrm{N}C \sqsubseteq D$ does not entail $\mathrm{N}C \equiv \mathrm{N}(C \sqcap D)$ (just as if C and $C \sqcap D$ were associated to different typicality orderings). Consequently, $\mathcal{DL}^{\mathrm{N}}$ does not satisfy the analogues of CT_T and CM_T either, that is, rules CT^N and CM^N in Table 4.

Example 8 Let $\mathcal{KB}$ be:

$$A \sqsubseteq_n E \quad (51) \qquad\qquad C \sqcap D \sqsubseteq_n \exists R.\neg E \quad (54)$$
$$C \sqsubseteq A \quad (52) \qquad\qquad C \sqcap D \sqsubseteq_n \forall R.N(C \sqcap D) \quad (55)$$
$$C \sqsubseteq_n D \quad (53)$$

[14]Strictly speaking, the conclusion of (50) does not satisfy the syntactic restrictions that the papers on typicality DLs pose on queries. Nonetheless, (50) is satisfied semantically; moreover, its conclusion can be equivalently expressed as a query. In particular, $T(C) \equiv T(C \sqcap D)$ is entailed by a knowledge base iff $(\neg T(C) \sqcup T(C \sqcap D)) \sqcap (T(C) \sqcup \neg T(C \sqcap D))(a)$ – where a is a fresh constant – is entailed.

Table 4: Candidate inference rules inspired by KLM postulates

Name	Rule schema	Sound in $\mathcal{DL}^{\mathrm{N}}$
REF^N	$\dfrac{}{\mathrm{N}C \sqsubseteq C}$	✓
CT^N	$\dfrac{\mathrm{N}C \sqsubseteq D \quad \mathrm{N}(C \sqcap D) \sqsubseteq E}{\mathrm{N}C \sqsubseteq E}$	partly
CM^N	$\dfrac{\mathrm{N}C \sqsubseteq D \quad \mathrm{N}C \sqsubseteq E}{\mathrm{N}(C \sqcap D) \sqsubseteq E}$	partly
LLE^N	$\dfrac{\mathrm{N}C \sqsubseteq E \quad C \equiv D}{\mathrm{N}D \sqsubseteq E}$	
RW^N	$\dfrac{\mathrm{N}C \sqsubseteq D \quad D \sqsubseteq E}{\mathrm{N}C \sqsubseteq E}$	✓
OR^N	$\dfrac{\mathrm{N}C \sqsubseteq E \quad \mathrm{N}D \sqsubseteq E}{\mathrm{N}(C \sqcup D) \sqsubseteq E}$	partly
RM^N	$\dfrac{\mathrm{N}C \sqsubseteq E \quad \mathrm{N}C \not\sqsubseteq \neg D}{\mathrm{N}(C \sqcap D) \sqsubseteq E}$	partly

C, D, and E range over $\mathcal{DL}$ concepts

The instances of $\mathrm{N}(C \sqcap D)$ satisfy the defeasible inclusions (54), (55), and (53), that is,

$$\begin{aligned}
\mathcal{KB} &\approx \mathrm{N}(C \sqcap D) \sqsubseteq \exists R.\neg E\,. \\
\mathcal{KB} &\approx \mathrm{N}(C \sqcap D) \sqsubseteq \forall R.N(C \sqcap D)\,, \\
\mathcal{KB} &\approx \mathrm{N}(C \sqcap D) \sqsubseteq D\,.
\end{aligned}$$

On the contrary, the lower-priority DI (51) is overridden, because the role restrictions of (54) and (55) imply that some instance x of $\mathrm{N}(C \sqcap D)$ satisfies $\neg E$, hence

$$\mathcal{KB} \not\approx \mathrm{N}(C \sqcap D) \sqsubseteq E\,. \tag{56}$$

The instances of $\mathrm{N}C$, instead, can consistently satisfy *all* the DIs in $\mathcal{KB}$, using the above x as a value of role R; it suffices to assume that x is *not* an instance of $\mathrm{N}C$. As explained above, there is no compelling reason for making x an instance of $\mathrm{N}C$, if we do not postulate that there exists a single, concept-independent normality criterion that forces $\mathrm{N}C$ to be equal to $\mathrm{N}(C \sqcap D)$. Consequently, the following entailments hold:

$$\mathcal{KB} \approx \mathrm{N}C \sqsubseteq D\,, \tag{57}$$

$$\mathcal{KB} \approx \mathrm{N}C \sqsubseteq E\,. \tag{58}$$

Note that (56), (57) and (58) show that CM^N is not satisfied in this example. There are two possible ways of satisfying CM^N here, that is, overriding (51) in $\mathrm{N}C$ (to remove (58)) or inferring $\mathrm{N}(C \sqcap D) \sqsubseteq E$. The former solution overrides a larger set of defeasible inclusions, in contrast with the principle that as many defaults as consistently possible should hold. Incidentally, this principle is embraced not only by Reiter's default logic, but also by Lehmann's approach to default reasoning based on the KLM postulates, cf. the presumptive reading of defaults and in particular Example 6 in [30], as well as Sec. 8 in the same paper. However, Lehmann's framework is propositional, so the above conflcit between maximal default application and CM could not be observed. The second solution makes $\mathrm{N}(C \sqcap D)$ inconsistent although there is no unresolved conflict in $\mathcal{KB}$. This solution would contrast with the principle that specificity settles conflicting defaults, since the inconsistency of $\mathrm{N}(C \sqcap D)$ can only be supported with (51), which is the least specific default in this example and should consequently be overridden by the other DIs, as it happens in $\mathcal{DL}^{\mathrm{N}}$.

A slight extension of the same example disproves CT^N. Extend $\mathcal{KB}$ with two more axioms:

$$A \sqsubseteq B \tag{59}$$

$$B \sqsubseteq_n \neg E\,. \tag{60}$$

Now, since axiom (51) is overridden in $\mathrm{N}(C \sqcap D)$, then (60) can be inherited and hence $\mathrm{N}(C \sqcap D) \sqsubseteq \neg E$. On the contrary, (60) is overridden by (51) in $\mathrm{N}C$, so $\mathrm{N}C \not\sqsubseteq \neg E$. This shows that CT^N does not hold. The above discussion on the consequences of CM^N in this example applies to CT^N as well.

The only other logics where a similar example can be modelled are typicality DLs, that can denote typical instances; this makes it possible to encode axiom (55). Typicality DLs satisfy CT_T and CM_T by overriding the analogue of (51). As we pointed out before, this is obtained through the assumption that the normality criterion is concept independent, which forces $\mathrm{N}C$ to collapse to $\mathrm{N}(C \sqcap D)$. Given the debatable consequences of this assumption (cf. the drawbacks of **and 1** and **or 2**), an interesting open question is whether weaker assumptions suffice to enforce internalized CT and CM. ∎

The next result shows that the occurrence of N in (55) is the only responsible for the falsification of CT^N and the dilemma of choosing between its unconditional application and the principle that a default should be applied unless higher priority or strong axioms state otherwise. More precisely, the next theorem proves that if $\mathcal{KB}$ is N-free (i.e. normality concepts are allowed only in queries), then the maximal default application principle itself spontaneously enforces CT^N, as well as two other postulates.[15]

Theorem 4 *If $\mathcal{KB}$ is N-free, then CT^N, OR^N, and RM^N are sound.*

An interesting observation is that the occurrences of N in $\mathcal{KB}$ are essential to assert that certain roles have normal values. So N-free $\mathcal{DL}^{\mathrm{N}}$ has the same expressiveness limitation as the DLs based on rational closure illustrated in Sec. 2.4, namely, defeasible axioms do not apply to role values. Given that removing this restriction in $\mathcal{DL}^{\mathrm{N}}$ affects some of the internalized postulates, an interesting open question is whether removing the same limitation from [20, 21, 22, 19, 16] may similarly interfere with the internalized postulates.

Differently from the above postulates, the soundness of CM^N is affected not only by the explicit occurrences of N, but also by the novel requirement that unresolved conflicts yield inconsistent prototypes.

Example 9 Consider the simple N-free $\mathcal{KB}$ consisting of:

$$D \sqsubseteq \neg E \quad (61)$$

$$C \sqsubseteq_n D \quad (62)$$

$$C \sqsubseteq_n E\,. \quad (63)$$

Here the defeasible inclusions (62) and (63) together contradict the strong inclusion (61), but they have the same priority and hence the conflict cannot be resolved in C, therefore its prototype is inconsistent ($\mathrm{N}C \sqsubseteq \bot$). On the contrary, the instances of $C \sqcap D$, by definition, satisfy D (be they normal or not), so (63) is overridden and does not belong to the prototype of $C \sqcap D$, which is consistent. This shows that CM^N does not hold. If it were applied, then $\mathrm{N}(C \sqcap D)$ would be made inconsistent, too, which is difficult to justify: it is not clear why the strong facts $\mathrm{N}(C \sqcap D) \sqsubseteq D$ and (61) should not override (63) in the prototype of $C \sqcap D$. ∎

There are no further reasons why CM^N is not valid, as shown by the following theorem:

Theorem 5 *If $\mathcal{KB}$ is N-free and $\mathrm{N}C$ is satisfiable w.r.t. $\mathcal{KB}$, then CM^N is sound.*

[15] The counterexamples to OR^N and RM^N are similar to Example 8. They are discussed in the appendix.

We leave two general questions open, namely, whether an overriding criterion that does not silently remove unresolved conflicts can be harmonized with the internalized version of CM, and whether this postulate should be restricted to consistent prototypes only. The typicality logic $\mathcal{ALC} + \mathbf{T}_{min}$ would deal with the above example by making C inconsistent, that is, the conflict between $T(C) \sqsubseteq D$ and $T(C) \sqsubseteq E$ (the analogues of (62) and (63)) is propagated from the prototype of C to all of C. In other words, if there can be no prototypical instances, there can be no atypical instances, either, while one would probably expect a clash between defeasible inclusions to affect typical instances only.

We are left to discuss REF^N, LLE^N, and RW^N. Rule REF^N is nothing but axiom (49), so it clearly holds without restrictions. Similarly, RW^N is unconditionally sound: it is a classical inference, and $\mathcal{DL}^{\mathrm{N}}$ is closed under classical reasoning. On the contrary, LLE^N is not sound, in general, although it should. In the "weak" semantics of $\mathcal{DL}^{\mathrm{N}}$ – where distinct normality concepts are treated like different concept names – $\mathrm{N}C$ and $\mathrm{N}D$ can be made artificially different (e.g. with axioms like $\top \sqsubseteq_n \exists R.(\mathrm{N}C \sqcap \neg \mathrm{N}D)$). In the N-free fragment (where such axioms cannot be asserted) we have slightly better properties:

Theorem 6 *If $\mathcal{KB}$ is N-free, then LLE^N is sound.*

Still, $C \equiv D$ does not imply $\mathrm{N}C \equiv \mathrm{N}D$ (LLE^N states only that $\mathrm{N}C$ and $\mathrm{N}D$ have the same logical properties), although this would be a reasonable inference to have, in order to enforce the natural principle that equivalents can be substituted for equivalents. In the next section, $\mathcal{DL}^{\mathrm{N}}$'s semantics is refined in order to satisfy the above equivalence and make LLE^N valid.

6 A refined semantics for $\mathcal{DL}^{\mathrm{N}}$

Hereafter, $\mathcal{DL}$ refers to any classical description logic that corresponds to a fragment of First Order Logic, such as $\mathcal{SROIQ}$ (the logic underlying OWL2).[16]

The new semantics is centred around a set of *bounded* interpretations that constitute a sort of Kripke structure, used to define overriding and constrain the semantics of normality concepts. We will see that, by the Löwenheim-Skolem theorem, this set of interpretations is representative of *all* interpretations.

Let Δ be an arbitrary but fixed infinite domain. A Δ-*interpretation* $\mathcal{I} = \langle \Delta^{\mathcal{I}}, \cdot^{\mathcal{I}} \rangle$ is any extension of a (classical) interpretation of $\mathcal{DL}$ such that $\Delta^{\mathcal{I}} \subseteq \Delta$ and $\mathrm{N}C^{\mathcal{I}} \subseteq C^{\mathcal{I}}$, for all concepts C in $\mathcal{DL}$. Given $\mathcal{KB} = \mathcal{S} \cup \mathcal{D}$, we say that $\mathcal{I}$ is a $(\mathcal{S}, \Delta)$-*premodel* iff (i) $\mathcal{I}$ is a classical model of $\mathcal{S}$ ($\mathcal{I} \models \mathcal{S}$), and (ii) $\mathrm{N}C^{\mathcal{I}} = \mathrm{N}D^{\mathcal{I}}$ if $C^{\mathcal{J}} = D^{\mathcal{J}}$ holds for all Δ-interpretation $\mathcal{J}$ satisfying $\mathcal{S}$. Premodels satisfy $\mathcal{S}$ and (49), as in the old semantics [5]; additionally, they make $\mathrm{N}C$ and $\mathrm{N}D$ equivalent whenever C and D are intensionally equivalent within the set of classical models of $\mathcal{S}$ bounded by Δ.

Recall that the set of normality concepts that *satisfy a DI* δ in a Δ-interpretation $\mathcal{I}$ is:

$$\mathsf{sat}(\delta, \mathcal{I}) = \{\mathrm{N}C \mid \forall x \in \mathrm{N}C^{\mathcal{I}},\ x \notin \mathsf{pre}(\delta)^{\mathcal{I}} \text{ or } x \in \mathsf{con}(\delta)^{\mathcal{I}}\}\,.$$

[16]Technically, we need $\mathcal{DL}$ to enjoy the Löwenheim-Skolem theorem.

As in [5], δ is *overridden in* $\mathrm{N}C$ *in the context of* $\mathcal{I}$ (for short, δ is overridden in $\mathrm{N}C/\mathcal{I}$) if δ makes the prototype of C necessarily inconsistent with the higher-priority axioms that are satisfied and not overridden in the context of $\mathcal{I}$ (the definition of overriding is recursive). The novelty is that by "necessarily", we mean that δ is inconsistent with those axioms in the *set* of all $(\mathcal{S},\Delta)$-premodels.[17] The set of $\mathrm{N}C$ such that δ is overridden in $\mathrm{N}C/\mathcal{I}$, denoted by $\mathsf{ovd}_{\mathcal{KB}}(\delta,\mathcal{I})$, is thus defined as follows:

Definition 1 (Overriding, ovd**)** *Let* $\mathcal{KB}=\mathcal{S}\cup\mathcal{D}$ *and* $\mathcal{I}$ *be a* Δ*-interpretation. Then,* $\mathrm{N}C\in\mathsf{ovd}_{\mathcal{KB}}(\delta,\mathcal{I})$ *(i.e.* δ *is overridden in* $\mathrm{N}C/\mathcal{I}$*) iff there exists no* $(\mathcal{S},\Delta)$*-premodel* $\mathcal{J}$ *satisfying the following conditions:*

1. $\mathrm{N}C\in\mathsf{sat}(\delta,\mathcal{J})$ *(*$\mathrm{N}C$ *satisfies* δ *in* $\mathcal{J}$*);*
2. $\mathrm{N}C^{\mathcal{J}}\neq\emptyset$ *(*$\mathrm{N}C$ *is consistent);*
3. *for all* $\delta'\in\mathcal{D}$*, if* $\delta'\prec\delta$*, then* $\mathsf{sat}(\delta',\mathcal{I})\setminus\mathsf{ovd}_{\mathcal{KB}}(\delta',\mathcal{I})\subseteq\mathsf{sat}(\delta',\mathcal{J})$ *(in* $\mathcal{J}$*, each normality concept satisfies at least all the non-overridden, higher-priority* δ' *that are satisfied in the context of* $\mathcal{I}$*).*

The subscript $\mathcal{KB}$ of ovd will be dropped whenever clear from context.

Now we are ready to define the semantics of $\mathcal{DL}^{\mathrm{N}}$. A Δ-interpretation $\mathcal{I}$ *satisfies an axiom* ϵ *in the context* Δ, in symbols $\mathcal{I}\approx_{\Delta}\epsilon$, if and only if:

1. ϵ is a concept inclusion, a role inclusion, or an assertion, and $\mathcal{I}\models\epsilon$;
2. ϵ is a DI and, for all normality concept NC, $NC\in\mathsf{sat}(\epsilon,\mathcal{I})\cup\mathsf{ovd}(\epsilon,\mathcal{I})$ (that is, ϵ is satisfied by all normality concepts where it is not overridden).

Then a $(\mathcal{S},\Delta)$-premodel $\mathcal{I}$ *is a* Δ*-model of* $\mathcal{KB}=\mathcal{S}\cup\mathcal{D}$, in symbols $\mathcal{I}\approx_{\Delta}\mathcal{KB}$, iff, for all $\delta\in\mathcal{D}$, $\mathcal{I}\approx_{\Delta}\delta$.

Finally, the notion of semantic consequence is defined as usual: an axiom ϵ is a Δ*-consequence* of $\mathcal{KB}$, written $\mathcal{KB}\approx_{\Delta}\epsilon$, iff for all Δ-models $\mathcal{I}$ of $\mathcal{KB}$, $\mathcal{I}\approx_{\Delta}\epsilon$.

The relativization to Δ of the new semantics prevents any form of *unrestricted comprehension* [27] in Definition 1 and in the definition of premodel. Relativization introduces no ambiguity, as proved in the following theorem.

Theorem 7 *Let* Δ_1 *and* Δ_2 *be two infinite sets. Then,* $\mathcal{KB}\approx_{\Delta_1}\epsilon$ *iff* $\mathcal{KB}\approx_{\Delta_2}\epsilon$*.*

Accordingly, we shall drop the subscript Δ from now on. The translation of $\mathcal{DL}^{\mathrm{N}}$ knowledge bases in classical $\mathcal{DL}$ needs one simple change: equation (4) should be replaced with

$$\mathcal{KB}_0^{\Sigma}=\mathcal{S}\cup\big\{\mathrm{N}C\sqsubseteq C\mid \mathrm{N}C\in\Sigma\big\}\cup\big\{\mathrm{N}C\equiv\mathrm{N}D\mid\mathcal{S}\models C\equiv D,\ \{\mathrm{N}C,\mathrm{N}D\}\subseteq\Sigma\big\}\,. \quad (64)$$

The new translation is correct, as stated by the following theorem. The expression $\mathcal{DL}^{\Sigma}$ denotes the extension of $\mathcal{DL}$ with the normality concepts in Σ.

[17] Since the set of $(\mathcal{S},\Delta)$-premodels is an analogue of a Kripke structure, overriding can be regarded as an implicit modal operator.

Theorem 8 *Let $\mathcal{KB}$ be a general $\mathcal{DL}^{\rm N}$ knowledge base, and let Σ be any finite set of normality concepts containing at least all the $\mathrm{N}C$ that occur in $\mathcal{KB}$. Let $\mathcal{KB}^\Sigma$ denote the new translation of $\mathcal{KB}$ where (4) is replaced by (64). For all subsumptions and assertions $\alpha \in \mathcal{DL}^\Sigma$,*

$$\mathcal{KB} \mid\!\approx \alpha \textit{ iff } \mathcal{KB}^\Sigma \models \alpha.$$

With this translation, all the complexity results of [5], all the results introduced in the previous sections, and all the examples and counterexamples illustrated so far can be immediately extended to the refined semantics.[18] Moreover, the additional property of premodels makes LLE^N sound even if some normality concepts occur in the DIs, because $\mathcal{KB} \mid\!\approx C \equiv D$ makes $\mathrm{N}C \equiv \mathrm{N}D$ true in all premodels.

Theorem 9 *For* all *canonical knowledge bases $\mathcal{KB}$, LLE^N is sound.*

Therefore, under the refined semantics, the restriction of $\mathcal{DL}^{\rm N}$ to N-free knowledge bases is almost rational with respect to the postulates in Table 4; the only difference is that CM^N applies only to consistent normality concepts, as discussed in the previous section. In general, CT^N, CM^N, OR^N, and RM^N are not valid, because $\mathcal{DL}^{\rm N}$ does not adopt a unique, concept-independent normality citerion, cf. Example 8 and Appendix C. The other postulates are unconditionally valid.

The examples illustrated in sections 2–5 are not affected by the change of semantics, because they involve no pair of equivalent normality concepts. The following example shows a case where the difference is visible, instead, and inferring the equivalence of two normality concepts plays an important role.

Example 10 Let $\mathcal{KB}$ be the knowledge base consisting of:

$$A \sqcap B \sqsubseteq_n C\,, \tag{65}$$

$$D \sqsubseteq_n \exists R.\mathrm{N}(A \sqcap B)\,. \tag{66}$$

The old semantics can infer both $\mathrm{N}(A \sqcap B) \sqsubseteq C$ and $\mathrm{N}(B \sqcap A) \sqsubseteq C$. However, it entails neither $\mathrm{N}(A \sqcap B) \equiv \mathrm{N}(B \sqcap A)$ nor $\mathrm{N}D \sqsubseteq \exists R.\mathrm{N}(B \sqcap A)$, although it does entail the similar inclusion $\mathrm{N}D \sqsubseteq \exists R.\mathrm{N}(A \sqcap B)$. On the contrary, both sentences are correctly entailed by the new semantics. ▮

7 Further related work

Nonmonotonic DLs have been obtained from virtually all nonmonotonic semantics [17, 3, 4, 23, 33, 9, 24, 25, 20, 22]. They have been extensively compared to $\mathcal{DL}^{\rm N}$ in terms of: conflict resolution and priority criteria; immunity to inheritance blocking and closed-world effects; computational complexity. The results are summarized in [5,

[18] One of the reviewers asked to include in the paper the following rather obvious remarks. The intrinsic complexity of deciding $\mid\!\approx$ and the complexity of computing $\mathcal{KB}^\Sigma$ are two different things. The results in [5] concern the former; they have been proved using Theorem 8 with the smallest possible Σ, that contains only the normality concepts occurring in $\mathcal{KB}$ and in α. If $\mathcal{KB}^\Sigma$ had to be computed for *arbitrary* Σ, then the computation time required for this purpose would obviously grow linearly with $|\Sigma|$.

Table 1], and show that $\mathcal{DL}^{\mathrm{N}}$ actually addresses successfully a number of issues related to knowledge engineering needs.

Propositional Default logic, Autoepistemic logic, and Circumscription have been given a proof-theoretic characterization by means of sequent calculi [10]. To the best of our knowledge, this approach has never been extended to DLs. Tableaux methods for some nonmonotonic DLs can be found, for example, in [26, 24].

Circumscribed DLs [9] are simply a syntactic variant of the traditional circumscription of first-order knowledge bases; so, using standard techniques similar to those introduced in [29], one can see that their consequence relation is preferential, in the metalevel sense captured by Table 1. Similarly, well-established arguments, similar to those used in [34], can be used to prove that Default DLs and Autoepistemic DLs [3, 4, 23] are not cumulative because their sceptical and credulous entailment relations do not satisfy CM, and their credulous entailment does not satisfy CT, either. On the other hand, Default and Autoepistemic DLs extend classical inference, so they satisfy REF, LLE, and RW. These observation and the results of Section 3 provide a direct comparison of the consequence relations of these DLs and $\mathcal{DL}^{\mathrm{N}}$.

The studies of the DLs based on rational closure, lexicographic closure, and typicality, such as [33, 24, 25, 20, 22], are centred around *internalized* versions of the KLM postulates, instead. Differently from $\mathcal{DL}^{\mathrm{N}}$, those logics use the same constructs for asserting default properties and for expressing their consequences, so there is only one internalized version of the postulates. Relative to the internalized postulates, the typicality logic $\mathcal{ALC} + \mathbf{T}_{min}$ is preferential, while the other logics cited above are rational; the probabilistic logic based on lexicographic entailment, however, satisfies OR relative to a nonstandard disjunction operator, cf. [33, Theorem 4.19].

Unfortunately, in these logics, the metalevel postulates have never been considered, which reduces the opportunities for comparing their consequence relations with $\approx$. Even if the aforementioned logics satisfy most or all of the internalized postulates, this does not mean that the metalevel postulates are satisfied, as shown below, although apparently the arguments provided by Kraus, Lehmann and Magidor in support of some metalevel postulates should apply verbatim.

In the DLs based on rational and lexicographic closures, every classical (strong) assertion $C \sqsubseteq D$ about the world can be encoded as a defeasible inclusion $C \sqcap \neg D \mathrel{|\!\sim} \bot$. In this form, the inclusion cannot be possibly overridden by any other axiom, due to the definition of rank, so it actually behaves as certain, non-defeasible information. As a special case, using classical equivalences, every strong axiom $B \sqsubseteq \bot$ can be equivalently expressed as $B \mathrel{|\!\sim} \bot$.

Now, Kraus, Lehmann and Magidor argue that given some defeasible background knowledge $\mathcal{KB}$ and two strong sentences α and β, if both $\mathcal{KB} \cup \{\alpha\}$ and $\mathcal{KB} \cup \{\beta\}$ entail a strong sentence γ, then also $\mathcal{KB} \cup \{\alpha \vee \gamma\}$ should entail γ (OR rule):

> *It is a valid principle of monotonic classical reasoning and does not imply monotonicity, therefore we tend to accept it* [...] *If we think that "if John attends the party then, normally, the evening will be great" and also that "if Cathy attends the party then, normally, the evening will be great" and hear that at least one of Cathy or John will attend the party, shouldn't we be tempted to join in?* [29]

Concerning the first sentence of this argument, we remark *en passant* that we do not see any compelling reasons for considering the maximization of the classical inferences that do not yield monotonicity a more important goal than addressing knowledge engineering needs. Still, whoever embraces the above motivations should be ready to accept the following argument. Let $\mathcal{KB}$ be

$$A_1 \mathrel{|\!\sim} B_1\,, \tag{67}$$

$$A_2 \mathrel{|\!\sim} B_2\,, \tag{68}$$

$$C \mathrel{|\!\sim} A_1 \sqcap A_2\,. \tag{69}$$

Moreover, let α and β, respectively, be $B_1 \mathrel{|\!\sim} \bot$ and $B_2 \mathrel{|\!\sim} \bot$, that are equivalent to the strong inclusions $B_1 \sqsubseteq \bot$ and $B_2 \sqsubseteq \bot$. It is not hard to see that, according to the rational closure defined in [20], both $\mathcal{KB} \cup \{\alpha\}$ and $\mathcal{KB} \cup \{\beta\}$ entail $C \mathrel{|\!\sim} \bot$, that is equivalent to the strong inclusion $C \sqsubseteq \bot$. Therefore, the argument in support of the OR rule says that $\mathcal{KB} \cup \{\alpha \vee \beta\}$ should entail that C is empty, as well. Note that $\alpha \vee \beta$ can be equivalently expressed using the universal role U as $\exists U.B_1 \sqcap \exists U.B_2 \mathrel{|\!\sim} \bot$.

However, rational closure yields a different result. The axiom $\alpha \vee \beta$ is too weak to entail the emptyness of any of A_1 and A_2, consequently C is given a higher rank than A_1 and A_2, and due to inheritance blocking, $\mathcal{KB} \cup \{\alpha \vee \beta\}$ does not entail that C is empty, thereby violating the metalevel OR rule.[19]

Additionally, some typicality DLs do not satisfy the original, metalevel version of RM; a counterexample can be found in Appendix F. Unfortunately, there is no complete picture of which metalevel postulates are satisfied, although the similarities between $\mathcal{ALC} + \mathbf{T}_{min}$ and circumscription make us conjecture that its consequence relation is preferential in the metalevel sense. Then, the internalized and metalevel properties of $\mathcal{ALC} + \mathbf{T}_{min}$ would perfectly match.

A few critical analyses of the KLM postulates and related properties can be found in the literature on typicality logics. In [24, Sec. 7.2], internalized RM is shown to be too strong in the context of the typicality logics extending $\mathcal{ALC} + \mathbf{T}$. In [12], it is proved that several postulates related to rational consequence relations and their representability in terms of ranked models cannot be simultaneously satisfied in the context of Booth's propositional typicality logic.

Recall that the DLs based on rational and lexicographic closures do not apply defeasible axioms to role values. This is probably due to the fact that the rational closure defined for the propositional framework is applied to conditionals $C \mathrel{|\!\sim} D$ as if C and D were propositional formulae. However, in first-order syntax, the intended meaning of such expressions corresponds in fact to a sort of open conditional like $C(x) \mathrel{|\!\sim} D(x)$. Such conditionals have been considered by Lehmann and Magidor in a short TARK'90 paper [31]. They felt the need for additional postulates for introducing and eliminating $\exists$. This predicate version of the KLM theory has not been fully developed. Resuming this line of work might be a key to overcoming the aforementioned limitation of the DLs based on rational and lexicographic closures.

[19] Strictly speaking, the rational closure of DLs has been defined for $\mathcal{ALC}$ only, while the above example makes use of the universal role, too. Nonetheless, this example is interesting because it shows one of the side effects of extending the rational closure construction to richer DLs.

Finally, we mention a preferential DL [15] that similarly to [25] adopts the syntax based on defeasible inclusions, like the rational closure of DLs, but the semantics is based on ranked interpretations like those of typicality logics. The emphasis is on inductive and abductive reasoning, which are captured by two dual defeasible inclusion operators, one of which is monotonic. The other operator is nonmonotonic and satisfies all the internalized postulates.

8 Conclusions

The semantics of N is almost completely determined by the application of defeasible inclusions and overriding. The underlying monotonic logic forces only the elementary inclusions $\mathrm{N}C \sqsubseteq C$. In this respect, $\mathcal{DL}^{\mathrm{N}}$ is more similar to rational closure than typicality logics, as the latter's monotonic fragment gives the typicality operator several nontrivial properties, such as (50) and the analogues of **and 1** and **or 2**. Note that (50) alone suffices to enforce CM_T and CT_T.

The application of defeasible inclusions and the overriding criterion of $\mathcal{DL}^{\mathrm{N}}$ (with its new conflict handling approach) make $\models\!\!\!\approx$ a cumulative consequence relation. The remaining rules, OR and RM, are sound under the additional assumption that consistent concepts have at least one normal instance. Interestingly, in typicality logics, the existence of normal individuals is forced by the monotonic layer of the semantics, through the smoothness property of the normality relation. Similarly, we plan to investigate variants of $\mathcal{DL}^{\mathrm{N}}$ where the existence assumption is cast into the monotonic part of the semantics. This should be done with some care, though: if too many individuals were forced to be normal, then the undesirable closed-world assumption effects described in [5], that affect typicality logics, might be introduced in $\mathcal{DL}^{\mathrm{N}}$.

The above results provide an immediate comparison with the consequence relations of Circumscribed DLs (that are preferential) and those of Default and Autoepistemic DLs (that are not cumulative). Unfortunately, not much is known about the metalevel postulates satisfied by the DLs based on typicality, rational closure, and lexicographic closure (although the arguments put forward by Kraus, Lehmann and Magidor to support some of their postulates are general enough to apply in these contexts as well). As of today, we only know that some typicality DLs and the rational closure of $\mathcal{ALC}$ extended with the universal role U violate RM and OR, respectively (cf. Section 7 and Appendix F), but this is not enough for an extended comparison at this level.

Next we studied the logical properties of N. With the goals of $\mathcal{DL}^{\mathrm{N}}$ in mind, we argued that the boolean operators should not be unconditionally forced to commute with N and that only the inclusion $\mathrm{N}\neg C \sqsubseteq \neg\mathrm{N}C$ should always hold. Concerning the internalized version of the KLM postulates listed in Table 4, the refined semantics of $\mathcal{DL}^{\mathrm{N}}$ unconditionally satisfies REF^N, LLE^N, and RW^N. Example 8, Example 9, and Appendix C show that, in general, the other postulates of Table 4 conflict with either the principle that as many defaults as consistently possible should be applied, or with the principle that conflicts should be resolved in favor of higher priority axioms (with strong inclusions as top priority axioms). Example 8 explains also that this clash between principles arises because $\mathcal{DL}^{\mathrm{N}}$ relativizes the notion of typicality, that is, it does *not* assume that such notion is unique and concept-independent. Moreover, as shown in Example 9, CM^N interferes with the novel conflict handling approach of $\mathcal{DL}^{\mathrm{N}}$. In particular, if $\mathrm{N}C$ is inconsistent, then CM^N forces all concepts $\mathrm{N}(C \sqcap D)$ to be inconsistent, too, thereby preventing any resolution of the conflicts in C's prototype, no matter which piece of strong information D is added to C (cf. Example 9). This observation raises two interesting questions: Should internalized cautious monotonicity hold universally in a logic where unresolved conflicts yield inconsistent prototypes? What is the right way of harmonizing the KLM postulates with the emerging knowledge engineering needs? These issues could not be observed nor discussed in the original KLM framework, because it adopted the traditional, silent conflict removal approach.

Typicality DLs preserve both the default application principle underlying default logic and the analogue of Table 4 (with the possible exception of RM) by means of strong axioms, based on the assumption that the notion of typicality is concept-independent (cf. Section 5 and Appendix B). This assumption has some debatable consequences (cf. example 6) and the authors of [25] plan to investigate logics with multiple, relativized typicality relations. However, as argued in Appendix B, the internalized postulates are so closely related to the concept-independence assumption that this objective is likely to be challenging. The issues related to the novel conflict resolution approach, instead, are not visible in typicality logics, because in general the conflicts arising from multiple inheritance are silently removed.[20]

The DLs based on rational closure recalled in Section 2.4 satisfy the analogue of Table 4 where each inclusion $\mathrm{N}C \sqsubseteq D$ corresponds to a conditional $C \mid\!\sim D$. In this case, the cost to be payed is the inability of applying defeasible axioms to role values (so our counterexamples cannot be modelled). Of course, $\mathcal{DL}^{\mathrm{N}}$ could be restricted in a similar way by allowing only N-free knowledge bases. It is interesting to note that, with this restriction, the internalized postulates satisfied by the two frameworks become almost identical. Using the refined semantics introduced in Sec. 6, the internalized versions of the KLM axioms illustrated in Table 4 are all satisfied with the partial exception of CM^N, that is guaranteed to apply only to *consistent* concepts $\mathrm{N}C$. As discussed above, we leave it as an open question whether CM^N should be applied also to the *inconsistent* normality concepts produced by the novel conflict handling approach of $\mathcal{DL}^{\mathrm{N}}$.

We have also reported a second internalized version of the KLM axioms, centred around defeasible inclusions (cf. Table 2), that has been technically studied in [5]. We enriched the illustration of these rules by highlighting their relationships with Lehmann's redefinition of default reasoning. Using the notation of [30], a DI $C \sqsubseteq_n D$ corresponds to a default $(C : D)$. Since defeasible inclusions have the same role as Reiter's defaults as well as Lehmann's, they have similar behavior: for all concepts $\mathrm{N}E$, a DI $C \sqsubseteq_n D$ is either overridden in $\mathrm{N}E$ (and hence irrelevant) or it is applied; in the latter case, it yields the same effects as the corresponding classical inclusion $\mathrm{N}E \sqcap C \sqsubseteq D$. Accordingly, all the rules of Table 2 hold when their premises are not overridden, analogously to Lehmann's approach.

Summarizing the comparison of the above logics, the metalevel KLM postulates are still partly unexplored, while their internalized versions are currently satisfied at a

[20]Only the inconsistency of C or a direct conflict such as $\{T(C) \sqsubseteq A, T(C) \sqsubseteq \neg A\}$ can make $T(C)$ inconsistent.

cost. In $\mathcal{DL}^{\mathrm{N}}$– whose goals are not completely aligned with those of the KLM framework – the dilemma is solved in favor of other desiderata, namely, full expressiveness (i.e. supporting role restriction to normal instances), relativized typicality, and tight correspondence between inconsistent prototypes and unresolved conflicts. It is not yet clear to what extent the various incarnations of the KLM postulates can be harmonized with the above desiderata, nor what the ideal tradeoff should be, for each of the possible applications of nonmonotonic reasoning listed in [35]. The answer to these questions probably lies in the study of first-order versions of the KLM theory supporting the novel conflict handling approach.

Should the KLM postulates turn out to be inherently incompatible with some of the other requirements, remarks 1 and 2 suggest an interesting direction for further research, namely, designing a flexible logic whose inferences can be tuned to application needs by including suitable axioms in the knowledge base.

Finally, note that the postulates studied here do not directly relate $\sqsubseteq_n$ and N. Currently, the mutual relations between DIs and their effects on normality concepts can be partially and indirectly evinced from the results applying to tables 2 and 4. A complete account requires studying hybrid versions of the postulates, e.g. those obtained by taking the premises from Table 2 and the consequents from Table 4. This investigation, that is needed to complete the analysis of $\mathcal{DL}^{\mathrm{N}}$, will be the subject of future work.

Acknowledgements

This research is currently being funded by the European Unions Horizon 2020 research and innovation programme under grant agreement N. 731601.

References

[1] A. Artale, D. Calvanese, R. Kontchakov, and M. Zakharyaschev. The DL-lite family and relations. *J. Artif. Intell. Res. (JAIR)*, 36:1–69, 2009.

[2] F. Baader, S. Brandt, and C. Lutz. Pushing the EL envelope. In *Proc. of the 19th Int. Joint Conf. on Artificial Intelligence, IJCAI-05*, pages 364–369. Professional Book Center, 2005.

[3] F. Baader and B. Hollunder. Embedding defaults into terminological knowledge representation formalisms. *J. Autom. Reasoning*, 14(1):149–180, 1995.

[4] F. Baader and B. Hollunder. Priorities on defaults with prerequisites, and their application in treating specificity in terminological default logic. *J. Autom. Reasoning*, 15(1):41–68, 1995.

[5] P. A. Bonatti, M. Faella, I. Petrova, and L. Sauro. A new semantics for overriding in description logics. *Artif. Intell.*, 222:1–48, 2015.

[6] P. A. Bonatti, M. Faella, and L. Sauro. EL with default attributes and overriding. In *Int. Semantic Web Conf. (ISWC 2010)*, volume 6496 of *LNCS*, pages 64–79. Springer, 2010.

[7] P. A. Bonatti, M. Faella, and L. Sauro. Adding default attributes to EL++. In W. Burgard and D. Roth, editors, *AAAI*. AAAI Press, 2011.

[8] P. A. Bonatti, M. Faella, and L. Sauro. Defeasible inclusions in low-complexity DLs. *J. Artif. Intell. Res. (JAIR)*, 42:719–764, 2011.

[9] P. A. Bonatti, C. Lutz, and F. Wolter. The complexity of circumscription in DLs. *J. Artif. Intell. Res. (JAIR)*, 35:717–773, 2009.

[10] P. A. Bonatti and N. Olivetti. Sequent calculi for propositional nonmonotonic logics. *ACM Trans. Comput. Log.*, 3(2):226–278, 2002.

[11] P. A. Bonatti, I. M. Petrova, and L. Sauro. Optimizing the computation of overriding. In *The Semantic Web - ISWC 2015 - 14th International Semantic Web Conference, Bethlehem, PA, USA, October 11-15, 2015, Proceedings, Part I*, volume 9366 of *Lecture Notes in Computer Science*, pages 356–372. Springer, 2015.

[12] R. Booth, G. Casini, T. A. Meyer, and I. J. Varzinczak. On the entailment problem for a logic of typicality. In Q. Yang and M. Wooldridge, editors, *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 2805–2811. AAAI Press, 2015.

[13] G. Brewka and J. Lang, editors. *Principles of Knowledge Representation and Reasoning: Proceedings of the Eleventh International Conference, KR 2008, Sydney, Australia, September 16-19, 2008*. AAAI Press, 2008.

[14] A. Britz, G. Casini, T. Meyer, and I. J. Varzinczak. Preferential role restrictions. In T. Eiter, B. Glimm, Y. Kazakov, and M. Krötzsch, editors, *Informal Proceedings of the 26th International Workshop on Description Logics, Ulm, Germany, July 23 - 26, 2013*, volume 1014 of *CEUR Workshop Proceedings*, pages 93–106. CEUR-WS.org, 2013.

[15] K. Britz, J. Heidema, and T. A. Meyer. Semantic preferential subsumption. In Brewka and Lang [13], pages 476–484.

[16] K. Britz, T. Meyer, and I. Varzinczak. Semantic foundation for preferential description logics. In D. Wang and M. Reynolds, editors, *Proc. of the 24th Australasian Joint Conference on Artificial Intelligence*, number 7106 in LNAI, pages 491–500. Springer-Verlag, 2011.

[17] M. Cadoli, F. Donini, and M. Schaerf. Closed world reasoning in hybrid systems. In *Proc. of ISMIS'90*, pages 474–481. Elsevier, 1990.

[18] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, and R. Rosati. DL-Lite: Tractable description logics for ontologies. In *Proc. of AAAI 2005*, pages 602–607, 2005.

[19] G. Casini, T. Meyer, K. Moodley, and R. Nortje. Relevant closure: A new form of defeasible reasoning for description logics. In E. Fermé and J. Leite, editors, *Logics in Artificial Intelligence - 14th European Conference, JELIA 2014, Funchal, Madeira, Portugal, September 24-26, 2014. Proceedings*, volume 8761 of *Lecture Notes in Computer Science*, pages 92–106. Springer, 2014.

[20] G. Casini and U. Straccia. Rational closure for defeasible description logics. In T. Janhunen and I. Niemelä, editors, *JELIA*, volume 6341 of *Lecture Notes in Computer Science*, pages 77–90. Springer, 2010.

[21] G. Casini and U. Straccia. Defeasible inheritance-based description logics. In T. Walsh, editor, *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, pages 813–818. IJCAI/AAAI, 2011.

[22] G. Casini and U. Straccia. Defeasible inheritance-based description logics. *J. Artif. Intell. Res. (JAIR)*, 48:415–473, 2013.

[23] F. M. Donini, D. Nardi, and R. Rosati. Description logics of minimal knowledge and negation as failure. *ACM Trans. Comput. Log.*, 3(2):177–225, 2002.

[24] L. Giordano, V. Gliozzi, N. Olivetti, and G. L. Pozzato. A non-monotonic description logic for reasoning about typicality. *Artif. Intell.*, 195:165–202, 2013.

[25] L. Giordano, V. Gliozzi, N. Olivetti, and G. L. Pozzato. Semantic characterization of rational closure: From propositional logic to description logics. *Artif. Intell.*, 226:1–33, 2015.

[26] S. Grimm and P. Hitzler. A preferential tableaux calculus for circumscriptive ALCO. In A. Polleres and T. Swift, editors, *RR*, volume 5837 of *Lecture Notes in Computer Science*, pages 40–54. Springer, 2009.

[27] T. Jech. *Set Theory*. Springer, 2006.

[28] Y. Kazakov. RIQ and SROIQ are harder than SHOIQ. In Brewka and Lang [13], pages 274–284.

[29] S. Kraus, D. J. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artif. Intell.*, 44(1-2):167–207, 1990.

[30] D. J. Lehmann. Another perspective on default reasoning. *Ann. Math. Artif. Intell.*, 15(1):61–82, 1995.

[31] D. J. Lehmann and M. Magidor. Preferential logics: the predicate calculus case. In R. Parikh, editor, *Proceedings of the 3rd Conference on Theoretical Aspects of Reasoning about Knowledge, Pacific Grove, CA, March 1990*, pages 57–72. Morgan Kaufmann, 1990.

[32] D. J. Lehmann and M. Magidor. What does a conditional knowledge base entail? *Artif. Intell.*, 55(1):1–60, 1992.

[33] T. Lukasiewicz. Expressive probabilistic description logics. *Artif. Intell.*, 172(6-7):852–883, 2008.

[34] D. Makinson. General theory of cumulative inference. In M. Reinfrank, J. de Kleer, M. L. Ginsberg, and E. Sandewall, editors, *Non-Monotonic Reasoning, 2nd International Workshop, Grassau, FRG, June 13-15, 1988, Proceedings*, volume 346 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 1988.

[35] J. McCarthy. Applications of circumscription to formalizing common sense knowledge. *Artificial Intelligence*, 28:89–116, 1986.

[36] A. L. Rector. Defaults, context, and knowledge: Alternatives for OWL-indexed knowledge bases. In *Pacific Symposium on Biocomputing*, pages 226–237. World Scientific, 2004.

[37] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.

[38] E. Sandewall. Defeasible inheritance with doubt index and its axiomatic characterization. *Artif. Intell.*, 174(18):1431–1459, 2010.

A Proofs for Section 3

Remark 3 *All the proofs and results of this section hold identically under both the old semantics [5] and the new semantics introduced in Section 6, since Theorem 8 is a perfect replacement for the previous translation correctness result.*

Theorem 1. *Rules CT, CM, and LLE are sound.*

Proof. *Soundness of CT and CM*. Let Σ be the set of normality concepts explicitly occurring in $\mathcal{KB} \cup \{\gamma\}$. By the correctness of the classical translation ([5, Theorem 1] and Theorem 8 for the old and new semantics, respectively), it suffices to prove

$$\mathcal{KB}^{\Sigma} \equiv (\mathcal{KB} \cup \{\alpha\})^{\Sigma} \tag{70}$$

assuming that the first premise of CT and CM holds:

$$\mathcal{KB} \approx \alpha\,. \tag{71}$$

By (71) and [5, Theorem 21], $\mathcal{S} \models \alpha$. Since $\mathcal{S}$ is included in all $\mathcal{KB}_i^{\Sigma}$, it follows, by a straightforward induction, that for all $i = 0, 1, \ldots, |\mathcal{D}|$,

$$(\mathcal{KB} \cup \{\alpha\})_i^{\Sigma} \equiv \mathcal{KB}_i^{\Sigma}\,. \tag{72}$$

This implies (70) by definition.

LLE. Similar: If $\models \alpha \equiv \beta$, then for all $i = 0, 1, \ldots, |\mathcal{D}|$,

$$(\mathcal{KB} \cup \{\alpha\})_i^{\Sigma} \equiv (\mathcal{KB} \cup \{\beta\})_i^{\Sigma}\,.$$

hence $(\mathcal{KB} \cup \{\alpha\})^{\Sigma} \equiv (\mathcal{KB} \cup \{\beta\})^{\Sigma}$. The details are left to the reader. ∎

Theorem 2. *Consider RM and the axiom γ occurring in it. Suppose that $\mathcal{KB}$ contains an axiom $\neg(\mathrm{N}C \sqsubseteq \bot)$ (or equivalent formulations, cf. Sec. 2.1) for each $\mathrm{N}C$ occurring in $\mathcal{KB} \cup \{\gamma\}$ such that C is satisfiable w.r.t. $\mathcal{KB}$. Then RM is sound.*

Proof. Let Σ be the set of normality concepts occurring in $\mathcal{KB}\cup\{\gamma\}$. By the correctness of the classical translation, it suffices to prove that

$$\mathcal{KB}^{\Sigma} \subseteq (\mathcal{KB} \cup \{\beta\})^{\Sigma} \tag{73}$$

under the theorem's hypothesis and the assumption that the second premise of RM ($\mathcal{KB} \not\approx \neg\beta$) holds, equivalently:

$$\mathcal{KB}^{\Sigma} \not\models \neg\beta\,. \tag{74}$$

Suppose (73) does not hold (we shall derive a contradiction). Then there must be $i > 0$ and $\mathrm{N}C \in \Sigma$ such that

$$\begin{aligned} \mathcal{KB}^{\Sigma}_{i-1} \downarrow_{\prec\delta_i} \cup\{\delta_i^{\mathrm{N}C}\} &\quad\not\models\quad \mathrm{N}C \sqsubseteq \bot & (75)\\ (\mathcal{KB}^{\Sigma}_{i-1} \downarrow_{\prec\delta_i}) \cup \{\beta\} \cup \{\delta_i^{\mathrm{N}C}\} &\quad\models\quad \mathrm{N}C \sqsubseteq \bot\,. & (76) \end{aligned}$$

By (75), we have two consequences: First, C must be satisfiable, otherwise $\mathrm{N}C \sqsubseteq \bot$. Second, $\delta_i^{\mathrm{N}C} \in \mathcal{KB}^{\Sigma}$, by construction of $\mathcal{KB}^{\Sigma}$. It follows, by (76), that $\mathcal{KB}^{\Sigma} \models \neg\beta \vee (\mathrm{N}C \sqsubseteq \bot)$. Consequently, by the hypothesis, $\mathcal{KB}^{\Sigma} \models \neg\beta$, which contradicts (74). This completes the soundness proof for RM. ∎

Theorem 3. *Consider OR and the axiom γ occurring in it. Suppose that $\mathcal{KB}$ contains an axiom $\neg(\mathrm{N}C \sqsubseteq \bot)$ (or equivalent formulations) for each $\mathrm{N}C$ occurring in $\mathcal{KB}\cup\{\gamma\}$ such that C is satisfiable w.r.t. $\mathcal{KB}$. Then OR is sound.*

Proof. The proof for OR and Theorem 2 are based on similar ideas. There are two possibilities:

a) for all $i > 0$ and $\mathrm{N}C \in \Sigma$, $\delta_i^{\mathrm{N}C} \in (\mathcal{KB} \cup \{\alpha\})^{\Sigma}$ iff $\delta_i^{\mathrm{N}C} \in (\mathcal{KB} \cup \{\beta\})^{\Sigma}$;

b) there exists $i > 0$ and $\mathrm{N}C \in \Sigma$ such that $\delta_i^{\mathrm{N}C}$ belongs to exactly one of $(\mathcal{KB} \cup \{\alpha\})^{\Sigma}$ and $(\mathcal{KB} \cup \{\beta\})^{\Sigma}$.

In case (a), it is not hard to see that $(\mathcal{KB} \cup \{\alpha\})^{\Sigma} \equiv (\mathcal{KB} \cup \{\alpha \vee \beta\})^{\Sigma} \cup \{\alpha\}$ and $(\mathcal{KB}\cup\{\beta\})^{\Sigma} \equiv (\mathcal{KB}\cup\{\alpha\vee\beta\})^{\Sigma}\cup\{\beta\}$ (by induction on the steps of the translations). The soundness of OR then follows from classical inferences (the details are left to the reader).

Next we focus on case (b). Assume without loss of generality that i is the least index that satisfies (b), and that $\delta_i^{\mathrm{N}C} \in (\mathcal{KB} \cup \{\alpha\})^{\Sigma} \setminus (\mathcal{KB} \cup \{\beta\})^{\Sigma}$ (the other case is symmetric). Accordingly, by definition,

$$\begin{aligned} (\mathcal{KB} \cup \{\alpha\})^{\Sigma}_{i-1} \downarrow_{\prec\delta_i} \cup\{\delta_i^{\mathrm{N}C}\} &\quad\not\models\quad \mathrm{N}C \sqsubseteq \bot & (77)\\ (\mathcal{KB} \cup \{\beta\})^{\Sigma}_{i-1} \downarrow_{\prec\delta_i} \cup\{\delta_i^{\mathrm{N}C}\} &\quad\models\quad \mathrm{N}C \sqsubseteq \bot & (78) \end{aligned}$$

(note that (77) implies that C is satisfiable, so $\mathcal{KB}$ contains an axiom equivalent to $\neg(\mathrm{N}C \sqsubseteq \bot)$). Moreover, by the minimality of i, it can be proved that

$$(\mathcal{KB} \cup \{\alpha\})^{\Sigma}_{i-1} \equiv (\mathcal{KB} \cup \{\alpha \vee \beta\})^{\Sigma}_{i-1} \cup \{\alpha\} \tag{79}$$

$$(\mathcal{KB} \cup \{\beta\})^{\Sigma}_{i-1} \equiv (\mathcal{KB} \cup \{\alpha \vee \beta\})^{\Sigma}_{i-1} \cup \{\beta\} . \tag{80}$$

From (78) and (80) it follows that $(\mathcal{KB} \cup \{\alpha \vee \beta\})^{\Sigma}_{i-1} \cup \{\delta_i^{\mathrm{N}C}\} \models \neg\beta \vee (\mathrm{N}C \sqsubseteq \bot)$. Morever, by (78) and (79), $\delta_i^{\mathrm{N}C} \in (\mathcal{KB} \cup \{\alpha \vee \beta\})^{\Sigma}$. It follows that $(\mathcal{KB} \cup \{\alpha \vee \beta\})^{\Sigma} \models \neg\beta \vee (\mathrm{N}C \sqsubseteq \bot)$, and hence, by hypothesis, $(\mathcal{KB} \cup \{\alpha \vee \beta\})^{\Sigma} \models \neg\beta$. Then, clearly, $(\mathcal{KB} \cup \{\alpha \vee \beta\})^{\Sigma} \models \alpha$ and $\mathcal{KB} \cup \{\alpha \vee \beta\} \approx\!\!\!\!\models \alpha$. By Theorem 1, the following instance of CT is sound:

$$\frac{\mathcal{KB} \cup \{\alpha \vee \beta\} \mathrel{|\!\!\approx} \alpha \quad \mathcal{KB} \cup \{\alpha \vee \beta\} \cup \{\alpha\} \mathrel{|\!\!\approx} \gamma}{\mathcal{KB} \cup \{\alpha \vee \beta\} \mathrel{|\!\!\approx} \gamma} .$$

We have already proved the first premise; the second premise is equivalent to the first premise of OR; the conclusion is the same as OR's. This proves the soundness of OR. ∎

Theorem 1. *Rules CT, CM, and LLE are sound.*

Proof. *Soundness of CT and CM*. Let Σ be the set of normality concepts explicitly occurring in $\mathcal{KB} \cup \{\gamma\}$. By the correctness of the classical translation ([5, Theorem 1] and Theorem 8 for the old and new semantics, respectively), it suffices to prove

$$\mathcal{KB}^{\Sigma} \equiv (\mathcal{KB} \cup \{\alpha\})^{\Sigma} \tag{81}$$

assuming that the first premise of CT and CM holds:

$$\mathcal{KB} \mathrel{|\!\!\approx} \alpha . \tag{82}$$

By (82) and [5, Theorem 21], $\mathcal{S} \models \alpha$. Since $\mathcal{S}$ is included in all $\mathcal{KB}^{\Sigma}_i$, it follows, by a straightforward induction, that for all $i = 0, 1, \ldots, |\mathcal{D}|$,

$$(\mathcal{KB} \cup \{\alpha\})^{\Sigma}_i \equiv \mathcal{KB}^{\Sigma}_i . \tag{83}$$

This implies (81) by definition.

LLE. Similar: If $\models \alpha \equiv \beta$, then for all $i = 0, 1, \ldots, |\mathcal{D}|$,

$$(\mathcal{KB} \cup \{\alpha\})^{\Sigma}_i \equiv (\mathcal{KB} \cup \{\beta\})^{\Sigma}_i .$$

hence $(\mathcal{KB} \cup \{\alpha\})^{\Sigma} \equiv (\mathcal{KB} \cup \{\beta\})^{\Sigma}$. The details are left to the reader. ∎

B Details for Section 5.1

Here we briefly provide some details about the tight relationships between the logical properties of T in typicality DLs and the assumption that the typicality relations $\leq^{\mathcal{I}}$ are concept-independent.

We first recall how **and 1** is satisfied thanks to this assumption. Recall that for all interpretations $\mathcal{I}$, and for *all* concepts C, $T(C)^{\mathcal{I}}$ is the set of $\leq^{\mathcal{I}}$-minimal members of $C^{\mathcal{I}}$, where $\leq^{\mathcal{I}}$ is the unique typicality relation associated to $\mathcal{I}$. Assume that x is an instance of $(T(C) \sqcap T(D))^{\mathcal{I}}$. We must prove that, in accordance with axiom **and 1**, x is also an instance of $(T(C \sqcap D))^{\mathcal{I}}$. Suppose not; we shall derive a contradiction. Clearly, $x \in (C \sqcap D)^{\mathcal{I}}$, by assumption, so it must be the case that x is not a $\leq^{\mathcal{I}}$*-minimal* such member. Then there exists $y <^{\mathcal{I}} x$ such that $y \in (C \sqcap D)^{\mathcal{I}}$. However, since this implies that $y \in C^{\mathcal{I}}$ and $y <^{\mathcal{I}} x$, it follows that x cannot be an instance of $T(C)^{\mathcal{I}}$ (which contradicts our initial assumption). Clearly, it is crucial for this argument that the same typicality relation be used to define $T(C)^{\mathcal{I}}$ and $T(C \sqcap D)^{I}$.

The second step consists in verifying that without the concept-independence assumption **and 1** would not hold, indeed. For this purpose, assume that $T(C)$, $T(D)$, and $T(C \sqcap D)$ are defined with three different relations, that we denote with $\leq^{\mathcal{I}}_{C}$, $\leq^{\mathcal{I}}_{D}$, and $\leq^{\mathcal{I}}_{C \sqcap D}$, respectively. We construct a counterexample $\mathcal{I}$ where $\Delta^{I} = \{a, b\}$, $a \leq^{\mathcal{I}}_{C} b$, $a \leq^{\mathcal{I}}_{D} b$, and $b \leq^{\mathcal{I}}_{C \sqcap D} a$. It can be immediately verified that $(T(C) \sqcap T(D))^{\mathcal{I}} = \{a\}$ while $(T(C \sqcap D))^{\mathcal{I}} = \{b\}$, so **and 1** is not satisfied. Note that all the orderings are both modular and smooth, so this counterexample is compatible with all the standard properties of the preferential and rational models of typicality.

With similar arguments, one can see that both **or 2** and (50) strictly depend on the concept-independence assumption; the details are straightforward, and left to the reader. Moreover, if (50) does not hold, then the same happens to CM_T and CT_T; it can be shown with a direct analogue of Example 8. It follows that if one is determined to satisfy the six postulates REF_T–RM_T without the drawbacks of the concept-independence assumption, then this assumption cannot be simply dropped; it must be carefully weakened.

C Counterexamples for Section 5.2

Remark 4 *All the examples of this section hold identically under both the old semantics [5] and the new semantics introduced in Section 6 since Theorem 8 is a perfect replacement for the previous translation correctness result.*

Note that Example 8 applies to RM^N, too. Rule RM^N has the same puzzling consequences as CM^N and CT^N that have been discussed in the example. The counterexample to unrestricted OR^N is based on similar ideas. $\mathcal{KB}$ is:

$$C \sqsubseteq_n \exists R.\neg E \quad (84) \qquad D \sqsubseteq_n \exists R.\neg E \quad (87)$$
$$C \sqsubseteq_n \forall R.\mathrm{N}(C \sqcup D) \quad (85) \qquad D \sqsubseteq_n \forall R.\mathrm{N}(C \sqcup D) \quad (88)$$
$$C \sqcup D \sqsubseteq A \quad (86) \qquad A \sqsubseteq_n E \quad (89)$$

The top-priority DIs (84)–(85) and (87)–(88) apply to $N(C \sqcup D)$, and produce the following consequences:

$$\begin{aligned} \mathrm{N}(C \sqcup D) &\sqsubseteq \exists R.\neg E \\ \mathrm{N}(C \sqcup D) &\sqsubseteq \forall R.\mathrm{N}(C \sqcup D)\,. \end{aligned}$$

These two inclusions (supported by strong axioms and top-priority DIs) are inconsistent with $\mathrm{N}(C \sqcup D) \sqsubseteq E$ and hence override (89) in $\mathrm{N}(C \sqcup D)$.

This does not happen with $\mathrm{N}C$ and $\mathrm{N}D$. They can consistently satisfy *all* the DIs of $\mathcal{KB}'$, therefore (89) entails the following inclusions:

$$\mathrm{N}C \sqsubseteq E \qquad\qquad \mathrm{N}D \sqsubseteq E\,.$$

Consequently, OR^N would infer $\mathrm{N}(C \sqcup D) \sqsubseteq E$, making $\mathrm{N}(C \sqcup D)$ unnecessarily inconsistent. Alternative solutions are: unnecessarily overriding either the top-priority axioms in $\mathrm{N}(C \sqcup D)$, or the DI (89) in $\mathrm{N}C$ and/or $\mathrm{N}D$.

D Proofs for Section 5

Remark 5 *All the proofs and results of this section hold identically under both the old semantics [5] and the new semantics introduced in Section 6 since Theorem 8 is a perfect replacement for the previous translation correctness result.*

The proofs for Section 5 use an interesting lemma that needs the following notation:

$$\mathrm{N}C \sim \mathrm{N}D \text{ iff for all DI } \delta \in \mathcal{KB},\ \delta^{\{\mathrm{N}C\}} \in \mathcal{KB}^{\{\mathrm{N}C\}} \Leftrightarrow \delta^{\{\mathrm{N}D\}} \in \mathcal{KB}^{\{\mathrm{N}D\}}\,,$$

Informally speaking, $\mathrm{N}C \sim \mathrm{N}D$ means that C and D have the same prototype in $\mathcal{KB}$. Furthermore,

$$\mathcal{S}[C/D]$$

denotes the result of substituting D for C in $\mathcal{S}$. The lemma states that if C subsumes D and $\mathrm{N}C$ is consistent with D, then C and D have the same prototype ($\mathrm{N}C \sim \mathrm{N}D$). This property is similar to RM, but in fact it can be used to prove the properties of all the other rules in Table 4.

Lemma 1 *If $\mathcal{KB} = \mathcal{S} \cup \mathcal{D}$ is N-free and $\mathcal{S} \models D \sqsubseteq C$, then either $\mathrm{N}C \sim \mathrm{N}D$ or $\mathcal{KB}^{\{\mathrm{N}C\}} \models \mathrm{N}C \sqsubseteq \neg D$.*

Proof. Suppose that $\mathrm{N}C \not\sim \mathrm{N}D$, and let i be the least index in the linearization of $\mathcal{D}$ such that $\delta_i^{\{\mathrm{N}C\}} \in \mathcal{KB}^{\{\mathrm{N}C\}} \not\Leftrightarrow \delta_i^{\{\mathrm{N}D\}} \in \mathcal{KB}^{\{\mathrm{N}D\}}$. We have to prove that $\mathcal{KB}^{\{\mathrm{N}C\}} \models \mathrm{N}C \sqsubseteq \neg D$. For $Y = C, D$, let $X^{\{\mathrm{N}Y\}}$ abbreviate $\mathcal{S}_{i-1}^{\{\mathrm{N}Y\}} \downarrow_{\prec \delta_i} \cup \{\delta_i^{\{\mathrm{N}Y\}}\}$, and recall that $\delta_i^{\{\mathrm{N}Y\}} \in \mathcal{KB}^{\{\mathrm{N}Y\}}$ iff $X^{\{\mathrm{N}Y\}} \not\models \mathrm{N}Y \sqsubseteq \bot$. By the minimality of i, and since (by hypothesis) $\mathcal{KB}$ is N-free and $\mathcal{S} \models D \sqsubseteq C$, we have

$$\begin{aligned} X^{\{\mathrm{N}D\}} &= X^{\{\mathrm{N}C\}}[\mathrm{N}C/\mathrm{N}D] \setminus \{\mathrm{N}D \sqsubseteq C\} \cup \{\mathrm{N}D \sqsubseteq D\} \\ &\equiv X^{\{\mathrm{N}C\}}[\mathrm{N}C/\mathrm{N}D] \cup \{\mathrm{N}D \sqsubseteq D\}\,. \end{aligned} \qquad (90)$$

Claim 1:

$$X^{\{\mathrm{N}D\}} \models \mathrm{N}D \sqsubseteq \bot \qquad (91)$$
$$X^{\{\mathrm{N}C\}} \not\models \mathrm{N}C \sqsubseteq \bot \qquad (92)$$

The claim is proved by contradiction: assume that $X^{\{\mathrm{N}C\}} \models \mathrm{N}C \sqsubseteq \bot$; then also $X^{\{\mathrm{N}C\}} \cup \{\mathrm{N}C \sqsubseteq D\} \models \mathrm{N}C \sqsubseteq \bot$; it follows by (90) and the fact that classical inference is insensitive to renaming, that $X^{\{\mathrm{N}D\}} \models \mathrm{N}D \sqsubseteq \bot$; but then $\delta_i^{\{\mathrm{N}C\}} \in \mathcal{KB}^{\{\mathrm{N}C\}} \Leftrightarrow \delta_i^{\{\mathrm{N}D\}} \in \mathcal{KB}^{\{\mathrm{N}D\}}$ (a contradiction). This proves Claim 1.

Now, (91) and (90) imply

$$X^{\{\mathrm{N}C\}}[\mathrm{N}C/\mathrm{N}D] \cup \{\mathrm{N}D \sqsubseteq D\} \models \mathrm{N}D \sqsubseteq \bot$$

and hence, equivalently, $X^{\{\mathrm{N}C\}}[\mathrm{N}C/\mathrm{N}D] \models \mathrm{N}D \sqsubseteq \neg D$. By renaming $\mathrm{N}D$ back to $\mathrm{N}C$ we get

$$X^{\{\mathrm{N}C\}} \models \mathrm{N}C \sqsubseteq \neg D\,.$$

Since $X^{\{\mathrm{N}C\}} \subseteq \mathcal{KB}^{\{\mathrm{N}C\}}$, the Lemma immediately follows. ∎

Since the following proofs use similar ideas, we provide full details only for CM^N.

Theorem 4. *If $\mathcal{KB}$ is N-free then CT^N, OR^N, and RM^N are sound.*

Proof. *Proof for CT^N.* Assume that the premises of CT^N hold, that is, $\mathcal{KB} \approx\!\!\!| \ \mathrm{N}C \sqsubseteq D$ and $\mathcal{KB} \approx\!\!\!| \ \mathrm{N}(C \sqcap D) \sqsubseteq E$. Equivalently,

$$\mathcal{KB}^{\{\mathrm{N}C\}} \models \mathrm{N}C \sqsubseteq D\,, \tag{93}$$

$$\mathcal{KB}^{\{\mathrm{N}(C\sqcap D)\}} \models \mathrm{N}(C \sqcap D) \sqsubseteq E\,. \tag{94}$$

We only have to prove that $\mathcal{KB}^{\{\mathrm{N}C\}} \models \mathrm{N}C \sqsubseteq E$. By Lemma 1, either $\mathrm{N}C \sim \mathrm{N}(C \sqcap D)$ or $\mathcal{KB}^{\{\mathrm{N}C\}} \models \mathrm{N}C \sqsubseteq \neg(C \sqcap D)$. The latter, by (93) and (49), implies that $\mathrm{N}C$ is inconsistent; then clearly $\mathcal{KB}^{\{\mathrm{N}C\}} \models \mathrm{N}C \sqsubseteq E$ and the theorem holds. We are left to prove it for the former case, that is, $\mathrm{N}C \sim \mathrm{N}(C \sqcap D)$, which implies (using the hypothesis that $\mathcal{KB}$ is N-free):

$$\begin{aligned} &\mathcal{KB}^{\{\mathrm{N}C\}} = \\ &= \mathcal{KB}^{\{\mathrm{N}(C\sqcap D)\}}[\mathrm{N}(C \sqcap D)/\mathrm{N}C] \setminus \{\mathrm{N}C \sqsubseteq C \sqcap D\} \cup \{\mathrm{N}C \sqsubseteq C\} \\ &\equiv \mathcal{KB}^{\{\mathrm{N}(C\sqcap D)\}}[\mathrm{N}(C \sqcap D)/\mathrm{N}C] \cup \{\mathrm{N}C \sqsubseteq C\} \qquad \text{by (93)}\,. \end{aligned}$$

Since classical inference is monotonic and insensitive to renaming, it follows from the above equivalence and (94) that $\mathcal{KB}^{\{\mathrm{N}C\}} \models \mathrm{N}C \sqsubseteq E$. This completes the proof for CT^N.

Proof for OR^N. Assume that the premises of OR^N hold, equivalently:

$$\mathcal{KB}^{\{\mathrm{N}C\}} \models \mathrm{N}C \sqsubseteq E\,, \tag{95}$$

$$\mathcal{KB}^{\{\mathrm{N}D\}} \models \mathrm{N}D \sqsubseteq E\,. \tag{96}$$

We shall prove $\mathcal{KB}^{\{\mathrm{N}(C\sqcup D)\}} \models \mathrm{N}(C \sqcup D) \sqsubseteq E$. By Lemma 1, applied to $C \sqsubseteq C \sqcup D$ and $D \sqsubseteq C \sqcup D$, there are four possibilities:

(i) $\mathrm{N}C \sim \mathrm{N}(C \sqcup D)$ and $\mathrm{N}D \sim \mathrm{N}(C \sqcup D)$;

(ii) $\mathrm{N}C \sim \mathrm{N}(C \sqcup D)$ and $\mathcal{KB}^{\{\mathrm{N}(C\sqcup D)\}} \models \mathrm{N}(C \sqcup D) \sqsubseteq \neg D$;

(iii) $\mathcal{KB}^{\{\mathrm{N}(C\sqcup D)\}} \models \mathrm{N}(C \sqcup D) \sqsubseteq \neg C$ and $\mathrm{N}D \sim \mathrm{N}(C \sqcup D)$;

(iv) $\mathcal{KB}^{\{\mathrm{N}(C\sqcup D)\}} \models \mathrm{N}(C \sqcup D) \sqsubseteq \neg C$ and $\mathcal{KB}^{\{\mathrm{N}(C\sqcup D)\}} \models \mathrm{N}(C \sqcup D) \sqsubseteq \neg D$.

Proof for case (i). As before, $\mathrm{N}C \sim \mathrm{N}(C \sqcup D)$, implies:

$$\mathcal{KB}^{\{\mathrm{N}C\}} \quad \equiv \quad \mathcal{KB}^{\{\mathrm{N}(C\sqcup D)\}}[\mathrm{N}(C \sqcup D)/\mathrm{N}C] \cup \{\mathrm{N}C \sqsubseteq C\}\,.$$

By this equivalence, (95), and renaming $\mathrm{N}C$ back to $\mathrm{N}(C \sqcup D)$:

$$\mathcal{KB}^{\{\mathrm{N}(C\sqcup D)\}} \cup \{\mathrm{N}(C \sqcup D) \sqsubseteq C\} \models \mathrm{N}(C \sqcup D) \sqsubseteq E\,. \tag{97}$$

Symmetrically, by $\mathrm{N}D \sim \mathrm{N}(C \sqcup D)$, we have:

$$\mathcal{KB}^{\{\mathrm{N}(C\sqcup D)\}} \cup \{\mathrm{N}(C \sqcup D) \sqsubseteq D\} \models \mathrm{N}(C \sqcup D) \sqsubseteq E\,. \tag{98}$$

By (97) and (98), $\mathcal{KB}^{\{\mathrm{N}(C\sqcup D)\}} \cup \{\mathrm{N}(C \sqcup D) \sqsubseteq C \sqcup D\} \models \mathrm{N}(C \sqcup D) \sqsubseteq E$; since $(\mathrm{N}(C \sqcup D) \sqsubseteq C \sqcup D) \in \mathcal{KB}^{\{\mathrm{N}(C\sqcup D)\}}$, the theorem holds in case (i).

Proof for case (ii). Since $\mathrm{N}(C \sqcup D) \sqsubseteq C \sqcup D$ holds (by axiom schema (49)) and $\mathcal{KB}^{\{\mathrm{N}(C\sqcup D)\}} \models \mathrm{N}(C \sqcup D) \sqsubseteq \neg D$ (by (ii)), we have

$$\mathcal{KB}^{\{\mathrm{N}(C\sqcup D)\}} \models \mathrm{N}(C \sqcup D) \sqsubseteq C\,.$$

Moreover, by $\mathrm{N}C \sim \mathrm{N}(C \sqcup D)$, (97) holds, as in case (i). It follows that $\mathcal{KB}^{\{\mathrm{N}(C\sqcup D)\}} \models \mathrm{N}(C \sqcup D) \sqsubseteq E$, which proves the theorem for case (ii).

Proof for case (iii). Symmetric to (ii).

Proof for case (iv). Recall that $\mathrm{N}(C \sqcup D) \sqsubseteq C \sqcup D$ holds by axiom schema (49), so (iv) implies that $\mathrm{N}(C \sqcup D)$ is inconsistent. As a consequence, $\mathcal{KB}^{\{\mathrm{N}(C\sqcup D)\}} \models \mathrm{N}(C \sqcup D) \sqsubseteq E$, so the theorem holds in case (iv), too. This completes the proof of OR^N's soundness.

Proof for RM^N. Suppose the premises of RM^N hold, equivalently:

$$\mathcal{KB}^{\{\mathrm{N}C\}} \quad \models \quad \mathrm{N}C \sqsubseteq E\,, \tag{99}$$

$$\mathcal{KB}^{\{\mathrm{N}C\}} \quad \not\models \quad \mathrm{N}C \sqsubseteq \neg D\,. \tag{100}$$

As before, (100), (49), and Lemma 1 (applied to the inclusion $C \sqcap D \sqsubseteq C$) entail $\mathrm{N}C \sim \mathrm{N}(C \sqcap D)$, which implies:

$$\mathcal{KB}^{\{\mathrm{N}(C\sqcap D)\}} \quad \equiv \quad \mathcal{KB}^{\{\mathrm{N}C\}}[\mathrm{N}C/\mathrm{N}(C \sqcap D)] \cup \{\mathrm{N}(C \sqcap D) \sqsubseteq C \sqcap D\}\,.$$

This equivalence and (99) imply $\mathcal{KB}^{\{\mathrm{N}(C\sqcap D)\}} \models \mathrm{N}(C \sqcap D) \sqsubseteq E$, therefore the conclusion of RM^N holds. ∎

Theorem 5. *If $\mathcal{KB}$ is N-free and $\mathrm{N}C$ is satisfiable w.r.t. $\mathcal{KB}$, then CM^N is sound.*

Proof. Assume the premises of CM^N hold, therefore $\mathcal{KB}^{\{\mathrm{N}C\}} \models \mathrm{N}C \sqsubseteq D$ and $\mathcal{KB}^{\{\mathrm{N}C\}} \models \mathrm{N}C \sqsubseteq E$. By Lemma 1 (applied to the inclusion $C \sqcap D \sqsubseteq C$), either $\mathrm{N}C \sim \mathrm{N}(C \sqcap D)$ or $\mathcal{KB}^{\{\mathrm{N}C\}} \models \mathrm{N}C \sqsubseteq \neg(C \sqcap D)$. The latter, by the first assumption and (49), implies that $\mathrm{N}C$ is inconsistent, contradicting the hypothesis. Then $\mathrm{N}C \sim \mathrm{N}(C \sqcap D)$. The rest of the proof is identical to the final part of RM^N's proof. ∎

Theorem 6. *If $\mathcal{KB}$ is N-free, then LLE^N is sound.*

Proof. Suppose that LLE^N's premises hold: $\mathcal{KB} \mathrel{|\!\!\approx} \mathrm{N}C \sqsubseteq E$ and $\mathcal{KB} \mathrel{|\!\!\approx} C \equiv D$. By [5, Theorem 21], the latter implies $\mathcal{S} \models C \equiv D$, ($\mathcal{S}$ is the strong part of $\mathcal{KB}$). By this entailment and N-freedom:

$$\begin{aligned} \mathcal{KB}^{\{\mathrm{N}D\}} &= \mathcal{KB}^{\{\mathrm{N}C\}}[\mathrm{N}C/\mathrm{N}D] \setminus \{\mathrm{N}D \sqsubseteq C\} \cup \{\mathrm{N}D \sqsubseteq D\} \\ &\equiv \mathcal{KB}^{\{\mathrm{N}C\}}[\mathrm{N}C/\mathrm{N}D]\,. \end{aligned} \tag{101}$$

The first assumption is equivalent to $\mathcal{KB}^{\{\mathrm{N}C\}} \models \mathrm{N}C \sqsubseteq E$; then, by (101), $\mathcal{KB}^{\{\mathrm{N}D\}} \models \mathrm{N}D \sqsubseteq E$ (since classical inferences are insensitive to renaming), which is equivalent to LLE^N's conclusion. ∎

E Proofs for Section 6

Throughout this section we adopt the refined definitions introduced in Sec. 6.

Theorem 7. *Let Δ_1 and Δ_2 be two infinite domains. Then, $\mathcal{KB} \mathrel{|\!\!\approx}_{\Delta_1} \epsilon$ iff $\mathcal{KB} \mathrel{|\!\!\approx}_{\Delta_2} \epsilon$.*

Proof. The thesis clearly holds if Δ_1 and Δ_2 have the same cardinality. So, w.l.o.g. we assume that $card(\Delta_1) < card(\Delta_2)$.

Claim 1: For each $(\mathcal{S}, \Delta_1)$-premodel $\mathcal{I}_1$ (resp. $(\mathcal{S}, \Delta_2)$-premodel $\mathcal{I}_2$) there exists $(\mathcal{S}, \Delta_2)$-premodel $\mathcal{I}_2$ (resp. $(\mathcal{S}, \Delta_1)$-premodel $\mathcal{I}_1$) such that

- for each classical axiom α, $\mathcal{I}_1 \models \alpha$ iff $\mathcal{I}_2 \models \alpha$;
- for each DI δ, $\mathsf{sat}(\delta, \mathcal{I}_1) = \mathsf{sat}(\delta, \mathcal{I}_2)$ and $\mathsf{ovd}(\delta, \mathcal{I}_1) = \mathsf{ovd}(\delta, \mathcal{I}_2)$.

Let $\mathcal{I}_1$ be a $(\mathcal{S}, \Delta_1)$-premodel. For the (upward) Löwenheim–Skolem theorem, there exists an Δ_2-interpretation $\mathcal{I}_2$ that satisfies the same FOL formulas as $\mathcal{I}_1$. Then, it is straightforward to see that $\mathcal{I}_2$ is a $(\mathcal{S}, \Delta_2)$-premodel and for each defeasible inclusion δ, $\mathsf{sat}(\delta, \mathcal{I}_1) = \mathsf{sat}(\delta, \mathcal{I}_2)$. By applying the (downward) Löwenheim–Skolem theorem, the converse can be proved similarly.

It remains to show by induction on the priority relation $\prec$ that for each DI δ, $\mathsf{ovd}(\delta, \mathcal{I}_1) = \mathsf{ovd}(\delta, \mathcal{I}_2)$. Base case, δ has maximal priority. Assume that $NC \in \mathsf{ovd}(\delta, \mathcal{I}_1)$ and let $\mathcal{J}_2$ be $(\mathcal{S}, \Delta_2)$-premodel. By the (downward) Löwenheim–Skolem theorem, there exists a Δ_1-interpretation $\mathcal{J}_1$ which satisfies the same FOL formulas as $\mathcal{J}_2$. Then, $\mathcal{J}_1$ is a $(\mathcal{S}, \Delta_1)$-premodel and, since $NC \in \mathsf{ovd}(\delta, \mathcal{I}_1)$, either $NC \notin \mathsf{sat}(\delta, \mathcal{J}_1)$ or $\mathcal{J}_1 \models NC \sqsubseteq \bot$. However, $\mathsf{sat}(\delta, \mathcal{J}_1) = \mathsf{sat}(\delta, \mathcal{J}_2)$ and $\mathcal{J}_1 \models NC \sqsubseteq \bot$ iff $\mathcal{J}_2 \models NC \sqsubseteq \bot$. Consequently, $NC \in \mathsf{ovd}(\delta, \mathcal{I}_2)$. Analogously, it can be proved that if $NC \in \mathsf{ovd}(\delta, \mathcal{I}_2)$, then $NC \in \mathsf{ovd}(\delta, \mathcal{I}_1)$.

Induction step. By hypothesis, for all $\delta' \in \mathcal{D}$, if $\delta' \prec \delta$, then $\mathsf{ovd}(\delta', \mathcal{I}_1) = \mathsf{ovd}(\delta', \mathcal{I}_2)$. Assume that $NC \in \mathsf{ovd}(\delta, \mathcal{I}_1)$, by definition for all $(\mathcal{S}, \Delta_1)$-premodels $\mathcal{J}$ one of the following holds: (i) $NC \notin \mathsf{sat}(\delta, \mathcal{J})$, (ii) $\mathcal{J} \models NC \sqsubseteq \perp$, or (iii) for some $\delta' \in \mathcal{D}$, $\delta' \prec \delta$ and $\mathsf{sat}(\delta', \mathcal{I}_1) \setminus \mathsf{ovd}_{\mathcal{KB}}(\delta', \mathcal{I}_1) \subseteq \mathsf{sat}(\delta', \mathcal{J})$. Now, let $\mathcal{J}_2$ be $(\mathcal{S}, \Delta_2)$-premodel. As before, by the Löwenheim–Skolem theorem, there exists a $(\mathcal{S}, \Delta_1)$-premodel such that $\mathsf{sat}(\delta, \mathcal{J}_1) = \mathsf{sat}(\delta, \mathcal{J}_2)$ and $\mathcal{J}_1 \models NC \sqsubseteq \perp$ iff $\mathcal{J}_2 \models NC \sqsubseteq \perp$. Since $\mathsf{sat}(\delta', \mathcal{I}_1) = \mathsf{sat}(\delta', \mathcal{I}_2)$ and by induction hypothesis $\mathsf{ovd}_{\mathcal{KB}}(\delta', \mathcal{I}_1) = \mathsf{ovd}_{\mathcal{KB}}(\delta', \mathcal{I}_2)$, it follows that $\mathcal{J}_2$ satisfies one of (i), (ii), or (iii) too. Subsequently, $NC \in \mathsf{ovd}(\delta, \mathcal{I}_2)$. The converse can be proved analogously.

Claim 2: $\mathcal{KB} \mathrel{|\!\!\approx}_{\Delta_1} \epsilon$ iff $\mathcal{KB} \mathrel{|\!\!\approx}_{\Delta_2} \epsilon$. Assume that $\mathcal{KB} \mathrel{|\!\!\approx}_{\Delta_1} \epsilon$ and let $\mathcal{I}_2$ be a model $\mathcal{KB}$ in Δ_2. By Claim 1, there exists a $(\mathcal{S}, \Delta_1)$-premodel $\mathcal{I}_1$ that satisfies the same classical axiom and defeasible inclusions as $\mathcal{I}_2$. Consequently, the following chain of facts hold: $\mathcal{I}_1$ is model of $\mathcal{KB}$ (by Claim 1); $\mathcal{I}_1 \mathrel{|\!\!\approx} \epsilon$ (by the assumption $\mathcal{KB} \mathrel{|\!\!\approx}_{\Delta_1} \epsilon$); $\mathcal{I}_2 \mathrel{|\!\!\approx} \epsilon$ (again, by Claim 1). Hence, $\mathcal{KB} \mathrel{|\!\!\approx}_{\Delta_2} \epsilon$. The converse can be proved analogously. ∎

Theorem 8. *Let $\mathcal{KB}$ be a general $\mathcal{DL}^{\mathrm{N}}$ knowledge base, and let Σ be any finite set of normality concepts containing at least all the* N*C that occur in $\mathcal{KB}$. Let $\mathcal{KB}^{\Sigma}$ denote the new translation of $\mathcal{KB}$ where (4) is replaced by (64). For all subsumptions and assertions $\alpha \in \mathcal{DL}^{\Sigma}$,*

$$\mathcal{KB} \mathrel{|\!\!\approx} \alpha \textit{ iff } \mathcal{KB}^{\Sigma} \models \alpha.$$

Proof. Similar to the proof of Theorem 1 in [5]. We just report the points where substantial changes are required.

Claim 2 of Lemma 24: If $\delta_i^{\mathrm{N}D} \in \mathcal{KB}^{\Sigma}$ and $\mathrm{N}D \in \mathsf{ovd}(\delta_i, \mathcal{I})$, then there exist $j < i$ and $\mathrm{N}E \in \Sigma$ such that $\delta_j \in \mathcal{KB}$, $\delta_j^{\mathrm{N}E} \notin \mathcal{KB}^{\Sigma}$, and $\mathrm{N}E \notin \mathsf{ovd}(\delta_j, \mathcal{I})$.

Assume the hypotheses of Claim 2 hold. The first one ($\delta_i^{ND} \in \mathcal{KB}^{\Sigma}$) implies (by definition of the sequence $\langle \mathcal{S}_i^{\Sigma} \rangle_i$) that $\mathcal{S}_{i-1}^{\Sigma} \downarrow_{\prec \delta_i} \cup \{\delta_i^{ND}\} \not\models \mathrm{N}D \sqsubseteq \perp$, therefore, there exists $\mathcal{J}_0$ satisfying $\mathcal{S}_{i-1}^{\Sigma} \downarrow_{\prec \delta_i}$ and $\delta_i^{\mathrm{N}D}$, such that $\mathrm{N}D^{\mathcal{J}_0} \neq \emptyset$. Let $\mathcal{J}$ be an extension of $\mathcal{J}_0$ such that, for all normality concepts $\mathrm{N}H \notin \Sigma$, we have that

$$\mathrm{N}H^{\mathcal{J}} = \begin{cases} \mathrm{N}F^{\mathcal{J}} & \text{if for some } \mathrm{N}F \in \Sigma,\ \mathcal{S} \models F \equiv H \\ \emptyset & \text{otherwise} \end{cases}$$

By construction, $\mathcal{J}$ agrees with $\mathcal{J}_0$ on the interpretation of all $\mathcal{DL}^{\Sigma}$ sentences. Since $\mathcal{J}_0$ satisfies $\mathcal{S}_0^{\Sigma}$, this means that $\mathcal{J}$ satisfies $\mathcal{S}$, $\mathrm{N}H \sqsubseteq H$ (for all $\mathrm{N}H \in \Sigma$), and $\mathrm{N}H \equiv \mathrm{N}H'$ (for all for all $\mathrm{N}H, \mathrm{N}H' \in \Sigma$ such that $\mathcal{S} \models H \equiv H'$).

Concerning the normality concepts $\mathrm{N}H \notin \Sigma$, we distinguish two cases. In the first case there exists $\mathrm{N}F \in \Sigma$ such that $\mathcal{S} \models F \equiv H$. By construction $\mathrm{N}H^{\mathcal{J}} = \mathrm{N}F^{\mathcal{J}}$ and, as seen above, $\mathcal{J}$ satisfies $\mathrm{N}F \sqsubseteq F$. Therefore, being $\mathcal{J}$ a model of $\mathcal{S}$, we have that $\mathcal{J}$ satisfies $\mathrm{N}H \sqsubseteq H$ as well. Moreover, for each G such that $\mathcal{S} \models G \equiv H$, by transitivity $\mathcal{S} \models G \equiv F$, consequently $\mathrm{N}H^{\mathcal{J}} = \mathrm{N}G^{\mathcal{J}}$. In the second case $\mathcal{S} \not\models F \equiv H$, for each $\mathrm{N}F \in \Sigma$. By construction, $\mathrm{N}H^{\mathcal{J}} = \emptyset$ and hence $\mathcal{J}$ trivially satisfies $\mathrm{N}H \sqsubseteq H$. Moreover, if $\mathcal{S} \models F \equiv H$, by hypothesis $\mathrm{N}F^{\mathcal{J}}$ is empty as well – otherwise $\mathrm{N}H$ would come under the previous case. Hence, $\mathcal{J}$ satisfies $\mathrm{N}H \equiv \mathrm{N}F$.

So, putting all this together, we have that $\mathcal{J}$ is a premodel of $\mathcal{S}$. Moreover, by construction, $\mathcal{J}$ satisfies conditions 1 and 2 of Def. 1.[21] If $\mathcal{J}$ satisfied also condition 3, then $\mathrm{N}D \notin \mathsf{ovd}(\delta_i, \mathcal{I})$, which contradicts the assumptions. Therefore $\mathcal{J}$ violates condition 3; it follows that there must be some $\delta_j \in \mathcal{KB}$ $(j < i)$ and some normality concept $\mathrm{N}E$ such that $\mathrm{N}E \in \mathsf{sat}(\delta_j, \mathcal{I}) \setminus \mathsf{ovd}(\delta_j, \mathcal{I})$ but

$$\mathrm{N}E \notin \mathsf{sat}(\delta_j, \mathcal{J})\,. \tag{102}$$

This fact has two consequences. First, if $\mathrm{N}E \notin \Sigma$, then $\mathrm{N}E^{\mathcal{J}} = \mathrm{N}F^{\mathcal{J}}$, for some $\mathrm{N}F \in \Sigma$, otherwise $\mathrm{N}E$ would vacuously satisfy δ_j, contradicting (102). Therefore, we can assume w.l.o.g. that $\mathrm{N}E \in \Sigma$. Second, since $\mathcal{J}$ and $\mathcal{J}_0$ agree on $\mathcal{DL}^{\Sigma}$ by construction, (102) implies $\mathcal{J}_0 \not\models \delta_j^{\mathrm{N}E}$ which is possible only if $\delta_j^{\mathrm{N}E} \notin \mathcal{KB}^{\Sigma}$. Finally, $\mathrm{N}E \notin \mathsf{ovd}(\delta_j, \mathcal{I})$ holds because $\mathrm{N}E \in \mathsf{sat}(\delta_j, \mathcal{I}) \setminus \mathsf{ovd}(\delta_j, \mathcal{I})$.

Lemma 25: Let $\mathcal{I}$ be a $\mathcal{DL}^{\mathrm{N}}$ model of $\mathcal{KB}$. For all normality concepts $\mathrm{N}C \in \Sigma$, and for all $\delta_i \in \mathcal{KB}$, $\mathrm{N}C \in \mathsf{ovd}(\delta_i, \mathcal{I})$ iff

$$\mathcal{S}^{\Sigma}_{i-1} \downarrow_{\prec \delta_i} \cup \{\delta_i^{\mathrm{N}C}\} \models \mathrm{N}C \sqsubseteq \bot\,.$$

The *if* direction does not need to be changed whereas the *only if* direction only requires to define the extension $\mathcal{J}$ of $\mathcal{J}_0$ as previously done in Claim 2 of Lemma 24. As proved above, $\mathcal{J}$ is a premodel of $\mathcal{S}$ and the proof then continues as it was. ∎

Theorem 9. *For all canonical knowledge bases $\mathcal{KB}$, LLE^N is sound.*

Proof. Suppose the premises of LLE^N hold:

$$\mathcal{KB} \;\approx\; \mathrm{N}C \sqsubseteq E \tag{103}$$

$$\mathcal{KB} \;\approx\; C \equiv D\,. \tag{104}$$

By [5, Theorem 21], (104) implies $\mathcal{S} \models C \equiv D$ ($\mathcal{S}$ is the strong part of $\mathcal{KB}$). Then, for any infinite Δ, all $(\mathcal{S}, \Delta)$-premodels $\mathcal{J}$ satisfy $\mathrm{N}C^{\mathcal{J}} \equiv \mathrm{N}D^{\mathcal{J}}$. Consequently, by definition of Δ-model and Δ-consequence, $\mathcal{KB} \approx \mathrm{N}C \equiv \mathrm{N}D$, and hence, by (103), $\mathcal{KB} \approx \mathrm{N}D \sqsubseteq E$ (the conclusion of LLE^N). ∎

F Typicality logic and metalevel RM

In this appendix we illustrate a counterexample to the metalevel version of RM in the typicality logic $\mathcal{ALC}+\mathbf{T}_{min}$.[22] This logic minimizes the sets of atypical instances in the models of the monotonic typicality logic $\mathcal{ALC}+\mathbf{T}$. Let $\mathcal{KB}$ be the following knowledge base:

$$T(A) \sqsubseteq \neg C \tag{105}$$

$$T(B_1) \sqsubseteq \neg C \tag{106}$$

$$T(B_2) \sqsubseteq \neg C \tag{107}$$

$$\top \sqsubseteq \exists R.\big((B_1 \sqcap C) \sqcup (B_2 \sqcap A \sqcap C)\big) \tag{108}$$

[21] Where $\mathrm{N}C$ and δ are replaced by $\mathrm{N}D$ and δ_i, respectively.

[22] We are grateful to Laura Giordano for her friendly feedback on this example, that helped in simplifying its presentation.

By (108) every model of $\mathcal{KB}$ in $\mathcal{ALC}$ +$\mathbf{T}_{min}$ contains one atypical individual x that belongs either to $B_1 \sqcap \neg A \sqcap \neg B_2$ or to $B_2 \sqcap A \sqcap \neg B_1$.

Due to the latter case, it follows that this $\mathcal{KB}$ does *not* entail $\neg(\top \sqsubseteq \forall R.A)$. However, it implies $A \sqcap B_1 \sqcap C \sqsubseteq \bot$. These are the effects of the minimization of atypical instances.

Now, RM says that $\mathcal{KB} \cup \{\top \sqsubseteq \forall R.A\}$ should still entail $A \sqcap B_1 \sqcap C \sqsubseteq \bot$. However, this is not the case. In the models of the extended knowledge base, x belongs either to $B_1 \sqcap A \sqcap \neg B_2$ or to $B_2 \sqcap A \sqcap \neg B_1$. The former case does not correspond to any model of the original $\mathcal{KB}$, and invalidates the conclusion $A \sqcap B_1 \sqcap C \sqsubseteq \bot$.

This example applies also to the nonmonotonic DL obtained by applying to the models of $\mathcal{ALC} + \mathbf{T_R}$ (that satisfy the internalized version of RM) the minimization of atypical instances adopted by $\mathcal{ALC}$ +$\mathbf{T}_{min}$.