

Relevant Closure: A New Form of Defeasible Reasoning for Description Logics

Giovanni Casini, Thomas Meyer, Kodylan Moodley, and Riku Nortjé

Centre for Artificial Intelligence Research (CSIR Meraka and UKZN), South Africa, email:
gcasini,tmeyer,kmoodley,rmortje@csir.co.za

Abstract. Among the various proposals for defeasible reasoning for description logics, Rational Closure, a procedure originally defined for propositional logic, turns out to have a number of desirable properties. Not only it is computationally feasible, but it can also be implemented using existing classical reasoners. One of its drawbacks is that it can be seen as too weak from the inferential point of view. To overcome this limitation we introduce in this paper two extensions of Rational Closure: Basic Relevant Closure and Minimal Relevant Closure. As the names suggest, both rely on defining a version of relevance. Our formalisation of relevance in this context is based on the notion of a justification (a minimal subset of sentences implying a given sentence). This is, to our knowledge, the first proposal for defining defeasibility in terms of justifications—a notion that is well-established in the area of ontology debugging. Both Basic and Minimal Relevant Closure increase the inferential power of Rational Closure, giving back intuitive conclusions that cannot be obtained from Rational Closure. We analyse the properties and present algorithms for both Basic and Minimal Relevant Closure, and provide experimental results for both Basic Relevant Closure and Minimal Relevant Closure, comparing it with Rational Closure.

1 Introduction

Description logics, or DLs [1], are central to many modern AI applications because they provide the logical foundations of formal ontologies. The past 20 years have witnessed many attempts to introduce defeasibility in a DL setting, ranging from preferential approaches [8, 9, 12, 18, 28] to circumscription [4–6, 30], amongst others [2, 15].

Preferential extensions of DLs based on the *KLM approach* [23, 25] are particularly promising for two reasons. Firstly, it provides a formal analysis of defeasible properties, which plays a central role in assessing how intuitive the obtained results are. And secondly, it allows for decision problems to be reduced to classical entailment checking, sometimes without blowing up the computational complexity with respect to the underlying classical case. The main disadvantage of the KLM approach is that its best known form of inferential closure, Rational Closure [25], can be seen as too weak from an inferential point of view. For example, it does not support the inheritance of defeasible properties. Suppose we know that mammalian and avian red blood cells are vertebrate red blood cells ($\text{MRBC} \sqsubseteq \text{VRBC}$, $\text{ARBC} \sqsubseteq \text{VRBC}$), that vertebrate red

blood cells normally have a cell membrane ($\text{VRBC} \sqsubseteq \exists \text{hasCM}.\top$), that vertebrate red blood cells normally have a nucleus ($\text{VRBC} \sqsubseteq \exists \text{hasN}.\top$), but that mammalian red blood cells normally don't ($\text{MRBC} \sqsubseteq \neg \exists \text{hasN}.\top$). Rational Closure allows us to conclude that avian vertebrate red blood cells normally have a cell membrane ($\text{ARBC} \sqsubseteq \exists \text{hasCM}.\top$), but not so for mammalian red blood cells ($\text{MRBC} \sqsubseteq \exists \text{hasCM}.\top$). Informally, the former can be concluded because avian red blood cells are a normal type of vertebrate red blood cell, while the latter can't because mammalian red blood cells are an abnormal type of vertebrate red blood cell.

In this paper we propose two new forms of defeasible reasoning to overcome this limitation. Both rely on the formalisation of a version of *relevance*. In resolving conflicts between sets of defeasible statements, we focus only on those that are *relevant* to the conflict, thereby ensuring that statements not involved in the conflict are guaranteed to be retained. For example, we regard $\text{VRBC} \sqsubseteq \exists \text{hasCM}.\top$ as *irrelevant* to the conflict between the three statements $\text{MRBC} \sqsubseteq \text{VRBC}$, $\text{VRBC} \sqsubseteq \exists \text{hasN}.\top$, and $\text{MRBC} \sqsubseteq \neg \exists \text{hasN}.\top$. As we shall see, this ensures that we can conclude, from both our new forms of defeasible reasoning, that $\text{MRBC} \sqsubseteq \exists \text{hasCM}.\top$.

The formal versions of relevance we employ are based on the notion of a *justification* – a minimal set of sentences responsible for a conflict [21]. We regard any sentence occurring in some justification as *potentially relevant* for resolving the conflict. All other sentences are deemed to be irrelevant to the conflict. Both Basic and Minimal Relevant Closure are based on the use of justifications. The difference between the two proposals is related to the way in which the relevant statements are chosen from among the potentially relevant ones.

Here we focus on the DL \mathcal{ALC} , although our definitions of Basic and Minimal Relevant Closure are applicable to any DL. The rest of the paper is structured as follows. First, we outline the DL \mathcal{ALC} and how it can be extended to represent defeasible information. Then we discuss existing approaches to defeasible reasoning for DLs, with a focus on Rational Closure. This is followed by presentations of our proposals for Basic Relevant Closure and Minimal Relevant Closure. We then consider the formal properties of our proposals, after which we present experimental results, comparing both Basic Relevant Closure and Minimal Relevant Closure with Rational Closure. Finally, we discuss related work and conclude with some indications of future work.

2 \mathcal{ALC} with Defeasible Subsumption

The language of the description logic \mathcal{ALC} is built up from a finite set of *concept names* $\mathbb{N}_{\mathcal{C}}$ and a finite set of *role names* $\mathbb{N}_{\mathcal{R}}$. The set of complex concepts (denoted by \mathcal{L}) is built in the usual way according to the rule:

$$C ::= A \mid \top \mid \perp \mid \neg C \mid C \sqcap C \mid C \sqcup C \mid \exists r.C \mid \forall r.C$$

The semantics of \mathcal{ALC} is the standard Tarskian semantics based on interpretations \mathcal{I} of the form $\mathcal{I} := \langle \Delta^{\mathcal{I}}, \cdot^{\mathcal{I}} \rangle$, where the domain $\Delta^{\mathcal{I}}$ is a non-empty set and $\cdot^{\mathcal{I}}$ is an *interpretation function* mapping concept names A in $\mathbb{N}_{\mathcal{C}}$ to subsets $A^{\mathcal{I}}$ of $\Delta^{\mathcal{I}}$ and role names r in $\mathbb{N}_{\mathcal{R}}$ to binary relations $r^{\mathcal{I}}$ over $\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$.

Given $C, D \in \mathcal{L}$, $C \sqsubseteq D$ is a (*classical*) *subsumption*. An \mathcal{ALC} *TBox* \mathcal{T} is a finite set of classical subsumptions. An interpretation \mathcal{I} *satisfies* $C \sqsubseteq D$ iff $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$.

Entailment of $C \sqsubseteq D$ by \mathcal{T} is defined in the standard (Tarskian) way. For more details on DLs the reader is referred to the Description Logic Handbook [1].

For \mathcal{ALC} with defeasible subsumption, or $\mathcal{ALC}(\sqsubseteq)$, we also allow *defeasible subsumptions* of the form $C \sqsubseteq D$, collected in a defeasible TBox, or DBox (a finite set of defeasible subsumptions). The semantics for $\mathcal{ALC}(\sqsubseteq)$ is obtained by augmenting every classical interpretation with an ordering on its domain [8, 18]. A ranked interpretation is a structure $\mathcal{R} = \langle \Delta^{\mathcal{R}}, \cdot^{\mathcal{R}}, \prec_{\mathcal{R}} \rangle$, where $\langle \Delta^{\mathcal{R}}, \cdot^{\mathcal{R}} \rangle$ is a DL interpretation and $\prec_{\mathcal{R}}$ is a modular ordering on $\Delta^{\mathcal{R}}$ satisfying the smoothness condition (for every $C \in \mathcal{L}$, if $C^{\mathcal{R}} \neq \emptyset$ then $\min_{\prec_{\mathcal{R}}}(C^{\mathcal{R}}) \neq \emptyset$), and where $\prec_{\mathcal{R}}$ is modular iff there is a ranking function $rk : X \rightarrow \mathbb{N}$ s.t. for every $x, y \in \Delta^{\mathcal{R}}$, $x \prec_{\mathcal{R}} y$ iff $rk(x) < rk(y)$. A defeasible subsumption $C \sqsubseteq D$ is satisfied in \mathcal{R} iff $\min_{\prec_{\mathcal{R}}}(C^{\mathcal{R}}) \subseteq D^{\mathcal{R}}$. Intuitively $C \sqsubseteq D$ is satisfied by \mathcal{R} whenever the *most normal* C s are also D s. It is easy to see that every ranked interpretation \mathcal{R} satisfies $C \sqsubseteq D$ iff \mathcal{R} satisfies $C \sqcap \neg D \sqsubseteq \perp$. That is, classical information can “masquerade” as defeasible information.

3 Reasoning with Defeasible Knowledge Bases

From a KR perspective it is important to obtain an appropriate form of defeasible entailment for $\mathcal{ALC}(\sqsubseteq)$. We shall deal with (defeasible) knowledge bases $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$, where \mathcal{T} is a (classical) finite TBox and \mathcal{D} a finite DBox. Given such a KB, the goal is to determine what (classical and defeasible) subsumption statements ought to follow from it. An obvious first attempt is to use the standard Tarskian notion of entailment applied to ranked interpretations: \mathcal{K} *preferentially entails* $C \sqsubseteq D$ iff every ranked interpretation satisfying all elements of \mathcal{K} also satisfies $C \sqsubseteq D$. However, it is known that this construction (known as Preferential Entailment) suffers from a number of drawbacks [25]. Firstly, it is *monotonic*—if $C \sqsubseteq D$ is in the Preferential Entailment of \mathcal{K} , then it is also in the Preferential Entailment of every $\mathcal{K}' = \langle \mathcal{T}', \mathcal{D}' \rangle$ such that $\mathcal{T} \subseteq \mathcal{T}'$ and $\mathcal{D} \subseteq \mathcal{D}'$. Secondly it is inferentially too weak—it does not support the inheritance of defeasible properties. An alternative to Preferential Entailment, first proposed by Lehmann et al. for the propositional case [25], and adapted to the DL case by Giordano et al. [17, 16] and Britz et al. [9], is that of *Rational Closure*. It is inferentially stronger than Preferential Entailment, is not monotonic, and has (limited) support for the inheritance of defeasible properties. An elegant semantic description of Rational Closure was recently provided by Giordano et al. for both the propositional case [19] and for $\mathcal{ALC}(\sqsubseteq)$ [17, 16]. Our focus here is on an algorithm for Rational Closure for $\mathcal{ALC}(\sqsubseteq)$, initially proposed by Casini and Straccia [12] and subsequently refined and implemented by Britz et al. [7]. A useful feature of the algorithm is that it reduces Rational Closure for $\mathcal{ALC}(\sqsubseteq)$ to classical entailment checking for \mathcal{ALC} . Below we define Rational Closure and present the algorithm.

$C \in \mathcal{L}$ is said to be *exceptional* for a knowledge base \mathcal{K} iff $\top \sqsubseteq \neg C$ is preferentially entailed by \mathcal{K} . Exceptionality checking can be reduced to classical entailment checking.

Proposition 1. *Britz et al. [7]: For a KB $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$, let $\overline{\mathcal{D}} = \{\neg D \sqcup E \mid D \sqsubseteq E \in \mathcal{D}\}$. For every $C \in \mathcal{L}$, $\top \sqsubseteq \neg C$ is preferentially entailed by \mathcal{K} iff $\mathcal{T} \models \prod \overline{\mathcal{D}} \sqsubseteq \neg C$.*

Exceptionality is used to build up a sequence of *exceptionality sets* E_0, E_1, \dots , and from this, an *exceptionality ranking* of concepts and defeasible subsumptions. Let

$\mathcal{E}_{\mathcal{T}}(\mathcal{D}) := \{C \sqsubseteq D \in \mathcal{D} \mid \mathcal{T} \models \bigwedge \overline{\mathcal{D}} \sqsubseteq \neg C\}$. Let $E_0 := \mathcal{D}$, and for $i > 0$, let $E_i := \mathcal{E}_{\mathcal{T}}(E_{i-1})$. It is easy to see that there is a smallest n such that $E_n = E_{n+1}$. The rank $r_{\mathcal{K}}(C)$ of $C \in \mathcal{L}$ is the smallest number r such that C is *not* exceptional for E_r . If C is exceptional for all E_i (for $i \geq 0$) then $r_{\mathcal{K}}(C) = \infty$. The rank $r_{\mathcal{K}}(C \sqsubseteq D)$ of any $C \sqsubseteq D$ is the rank $r_{\mathcal{K}}(C)$ of its antecedent C .

Definition 1. *Lehmann et al. [25], Britz et al. [9]:* $C \sqsubseteq D$ is in the Rational Closure of \mathcal{K} iff $r_{\mathcal{K}}(C) < r_{\mathcal{K}}(C \sqcap \neg D)$ or $r_{\mathcal{K}}(C) = \infty$.

Having defeasible subsumptions with infinite rank in the DBox is problematic from an algorithmic point of view because it does not allow for a clear separation of classical information (in the TBox \mathcal{T}) and defeasible information (in the DBox \mathcal{D}) in a knowledge base \mathcal{K} .

Definition 2. A knowledge base $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$ is well-separated iff $r_{\mathcal{K}}(C \sqsubseteq D) \neq \infty$ for every $C \sqsubseteq D \in \mathcal{D}$.

We will frequently assume knowledge bases to be well-separated. It is worth pointing out that this assumption is not a restriction of any kind, since every knowledge base can be converted into a well-separated one, as shown by Britz et al. [7].

Below we present a high-level version of the algorithm for Rational Closure implemented by Casini et al. [11]. It takes as input a *well-separated* KB $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$ and a query $C \sqsubseteq D$, and returns **true** iff the query is in the Rational Closure of \mathcal{K} . It also assumes the existence of a *partition* procedure which computes the ranks of the subsumptions in \mathcal{D} and partitions \mathcal{D} into n equivalence classes according to rank: $i = 0, \dots, n$, $\mathcal{D}_i := \{C \sqsubseteq D \mid r_{\mathcal{K}}(C) = i\}$. Note that, because \mathcal{K} is well-separated, none of the elements of \mathcal{D} will have infinite rank. The *partition* procedure performs at most a polynomial number of classical entailment checks to compute the ranks. The remaining part of the algorithm performs a linear number of classical entailment checks (in the size of \mathcal{D}).

Algorithm 1: Rational Closure

Input: A well-separated KB $\langle \mathcal{T}, \mathcal{D} \rangle$ and a query $C \sqsubseteq D$
Output: **true** iff $C \sqsubseteq D$ is in the Rational Closure of $\langle \mathcal{T}, \mathcal{D} \rangle$

- 1 $(\mathcal{D}_0, \dots, \mathcal{D}_n, n) := \text{partition}(\mathcal{D})$;
- 2 $i := 0$; $\mathcal{D}' := \mathcal{D}$;
- 3 **while** $\mathcal{T} \models \bigwedge \overline{\mathcal{D}'} \sqsubseteq \neg C$ **and** $\mathcal{D}' \neq \emptyset$ **do**
- 4 $\mathcal{D}' := \mathcal{D}' \setminus \mathcal{D}_i$; $i := i + 1$;
- 5 **return** $\mathcal{T} \models \overline{\mathcal{D}'} \sqcap C \sqsubseteq D$;

Informally, the algorithm keeps on removing defeasible subsumptions from \mathcal{D} , starting with the lowest rank, and proceeding rank by rank, until it finds the first DBox \mathcal{D}' for which C is no longer exceptional. $C \sqsubseteq D$ is then taken to be in the Rational Closure of \mathcal{K} iff $\mathcal{T} \models \overline{\mathcal{D}'} \sqcap C \sqsubseteq D$. Observe that, since every classical subsumption $C \sqsubseteq D$ can be rewritten as a defeasible subsumption $C \sqcap \neg D \sqsubseteq \perp$, Algorithm 1 is, indirectly,

able to deal with classical queries (of the form $C \sqsubseteq D$) as well. The same holds for the other algorithms defined in this paper.

To see how the algorithm works, consider the following example, which we use as a running example in the rest of the paper.

Example 1. We know that both avian red blood cells and mammalian red blood cells are vertebrate red blood cells, and that vertebrate red blood cells normally have a cell membrane. We also know that vertebrate red blood cells normally have a nucleus, but that mammalian red blood cells normally don't. We can represent this information in the KB $\mathcal{K}^1 = \langle \mathcal{T}^1, \mathcal{D}^1 \rangle$ with $\mathcal{T}^1 = \{\text{ARBC} \sqsubseteq \text{VRBC}, \text{MRBC} \sqsubseteq \text{VRBC}\}$ and $\mathcal{D}^1 = \{\text{VRBC} \sqsubseteq \exists \text{hasCM}.\top, \text{VRBC} \sqsubseteq \exists \text{hasN}.\top, \text{MRBC} \sqsubseteq \neg \exists \text{hasN}.\top\}$.

We get $\mathcal{D}_0^1 = \{\text{VRBC} \sqsubseteq \exists \text{hasN}.\top, \text{VRBC} \sqsubseteq \exists \text{hasCM}.\top\}$, and $\mathcal{D}_1^1 = \{\text{MRBC} \sqsubseteq \neg \exists \text{hasN}.\top\}$. Given the query $\text{ARBC} \sqsubseteq \text{hasCM}.\top$, $\mathcal{D}' = \mathcal{D}^1$ in line 5, from which it follows that the query is in the Rational Closure of \mathcal{K}^1 . Given the query $\text{MRBC} \sqsubseteq \text{hasCM}.\top$, however, we get $\mathcal{D}' = \mathcal{D}_1^1$ in line 5, and so this query is not in the Rational Closure of \mathcal{K}^1 . An analysis of the latter query turns out to be very instructive for our purposes here. Observe that, to obtain \mathcal{D}' , the algorithm removes all elements of $\mathcal{D}_0^1 = \{\text{VRBC} \sqsubseteq \exists \text{hasN}.\top, \text{VRBC} \sqsubseteq \exists \text{hasCM}.\top\}$ from \mathcal{D}^1 . Informally, the motivation for the removal of $\text{VRBC} \sqsubseteq \exists \text{hasN}.\top$ is easy to explain: together with $\text{MRBC} \sqsubseteq \text{VRBC}$ and $\text{MRBC} \sqsubseteq \neg \exists \text{hasN}.\top$ it is responsible for MRBC being exceptional. It is less clear, intuitively, why the defeasible subsumption $\text{VRBC} \sqsubseteq \exists \text{hasCM}.\top$ has to be removed. One could make the case that since it plays no part in the exceptionality of MRBC , it should be retained. As we shall discuss in the next section, this argument forms the basis of an approach to defeasible reasoning based on the *relevance* of defeasible subsumptions.

4 Relevant Closure

Here we outline our proposal for a version of defeasible reasoning based on relevance. The principle is an obvious abstraction of the argument outlined at the end of the previous section—identify those defeasible subsumptions deemed to be *relevant* w.r.t. a given query, and consider only these ones as being eligible for removal during the execution of the Rational Closure algorithm. More precisely, suppose we have identified $R \subseteq \mathcal{D}$ as the defeasible subsumptions relevant to the query $C \sqsubseteq D$. First we ensure that all elements of \mathcal{D} that are *not* relevant to the query are *not* eligible for removal during execution of the Rational Closure algorithm. For $R \subseteq \mathcal{D}$ let $\text{Rel}^{\mathcal{K}}(R) := \langle R, R^- \rangle$, where $R^- = \mathcal{D} \setminus R$. That is, R^- is the set of all the defeasible subsumptions that are not eligible for removal since they are not relevant w.r.t. the query $C \sqsubseteq D$. Then we apply a variant of Algorithm 1 (the Rational Closure algorithm) to \mathcal{K} in which the elements of R^- are not allowed to be eliminated. The basic algorithm for *Relevant Closure* is outlined below (Algorithm 2). Note that, as in the case of Algorithm 1, we assume that the knowledge base is well-separated. We say that a defeasible subsumption $C \sqsubseteq D$ is in the *Relevant Closure* of (a well-separated) \mathcal{K} w.r.t. a set of relevant defeasible subsumptions R iff the Relevant Closure algorithm (Algorithm 2) returns **true**, with \mathcal{K} , $C \sqsubseteq D$, and $\text{Rel}^{\mathcal{K}}(R)$ as input.

For Example 1, an appropriate choice for R would be the set $\{\text{VRBC} \sqsubseteq \exists \text{hasN}.\top, \text{MRBC} \sqsubseteq \neg \exists \text{hasN}.\top\}$ since these are the two defeasible subsumptions responsible for

Algorithm 2: Relevant Closure

Input: A well-separated KB $\langle \mathcal{T}, \mathcal{D} \rangle$, a query $C \sqsubseteq D$, and the partition

$$Rel^{\mathcal{K}}(R) = \langle R, R^- \rangle$$

Output: **true** iff $C \sqsubseteq D$ is in the Relevant Closure of $\langle \mathcal{T}, \mathcal{D} \rangle$

- 1 $(\mathcal{D}_0, \dots, \mathcal{D}_n, n) := \text{partition}(\mathcal{D})$;
 - 2 $i := 0$; $R' := R$;
 - 3 **while** $\mathcal{T} \models \bigcap \overline{R^-} \cap \bigcap \overline{R'} \sqsubseteq \neg C$ **and** $R' \neq \emptyset$ **do**
 - 4 $R' := R' \setminus (\mathcal{D}_i \cap R)$; $i := i + 1$;
 - 5 **return** $\mathcal{T} \models \bigcap \overline{R^-} \cap \bigcap \overline{R'} \cap C \sqsubseteq D$;
-

MRBC being exceptional (w.r.t. \mathcal{K}). If $R = \{\text{VRBC} \sqsubseteq \exists \text{hasN.T}, \text{MRBC} \sqsubseteq \neg \exists \text{hasN.T}\}$ and $R^- = \{\text{VRBC} \sqsubseteq \exists \text{hasCM.T}\}$ (that is, it is information not eligible for removal), it is easy to see that we can derive $\text{MRBC} \sqsubseteq \text{hasCM.T}$, since $\{\text{ARBC} \sqsubseteq \text{VRBC}, \text{MRBC} \sqsubseteq \text{VRBC}\} \models (\neg \text{VRBC} \sqcup \text{hasCM.T}) \cap (\neg \text{MRBC} \sqcup \neg \exists \text{hasN.T}) \cap \text{MRBC} \sqsubseteq \text{hasCM.T}$.

4.1 Basic Relevant Closure

The explanation above still leaves open the question of how to define relevance w.r.t. a query. The key insight in doing so, is to *associate relevance with the subsumptions responsible for making the antecedent of a query exceptional*. We shall refer to such sets of subsumptions as *justifications*.

Definition 3. For $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$, $\mathcal{J} \subseteq \mathcal{D}$, and $C \in \mathcal{L}$, \mathcal{J} is a C -justification w.r.t. \mathcal{K} iff C is exceptional for $\langle \mathcal{T}, \mathcal{J} \rangle$ (i.e. $\top \sqsubseteq \neg C$ is in the Preferential Entailment of $\langle \mathcal{T}, \mathcal{J} \rangle$) and for every $\mathcal{J}' \subset \mathcal{J}$, C is not exceptional for \mathcal{J}' .

The choice of the term *justification* is not accidental, since it closely mirrors the notion of a justification for classical DLs, where a justification for a sentence α is a minimal set implying α [21]; it corresponds to the notion of *kernel*, used a lot in base-revision literature [20]. Given the correspondence between exceptionality and classical entailment in Proposition 1, the link is even closer.

Corollary 1. \mathcal{J} is a C -justification w.r.t. $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$ iff $\mathcal{J} \subseteq \mathcal{D}$, $\mathcal{T} \models \overline{\mathcal{J}} \sqsubseteq \neg C$, and for every $\mathcal{J}' \subset \mathcal{J}$, $\mathcal{T} \not\models \overline{\mathcal{J}'} \sqsubseteq \neg C$.

This places us in a position to define our first relevance-based version of defeasible reasoning. We identify *relevance* for a query $C \sqsubseteq D$ with all subsumptions occurring in some C -justification for \mathcal{K} . For $C \in \mathcal{L}$, and a KB \mathcal{K} , let $\mathcal{J}^{\mathcal{K}}(C) = \{\mathcal{J} \mid \mathcal{J} \text{ is a } C\text{-justification w.r.t. } \mathcal{K}\}$.

Definition 4. $C \sqsubseteq D$ is in the Basic Relevant Closure of \mathcal{K} iff it is in the Relevant Closure of \mathcal{K} w.r.t. $\bigcup \mathcal{J}^{\mathcal{K}}(C)$.

For Example 1, $\mathcal{J} = \{\text{VRBC} \sqsubseteq \exists \text{hasN.T}, \text{MRBC} \sqsubseteq \neg \exists \text{hasN.T}\}$ is the one MRBC-justification for D^1 , and $\text{MRBC} \sqsubseteq \text{hasCM.T}$ is in the Basic Relevant Closure of \mathcal{K}^1 as seen above, since the axiom $\text{VRBC} \sqsubseteq \exists \text{hasCM.T}$ is not in any MRBC-justification and is therefore deemed to be irrelevant w.r.t. the query.

To summarise, unlike Rational Closure, Basic Relevant Closure ensures that the de-feasible property of having a cell membrane is inherited by mammalian red blood cells from vertebrate red blood cells, even though mammalian red blood cells are abnormal vertebrate red blood cells (in the sense of not having a nucleus).

4.2 Minimal Relevant Closure

Although Basic Relevant Closure is inferentially stronger than Rational Closure, it can still be viewed as inferentially too weak, since it views *all* subsumptions occurring in some C -justifications as relevant, and therefore eligible for removal. In particular, it does not make proper use of the *ranks* of the subsumptions in a DBox. In this section we strengthen the notion of relevance by identifying it with the subsumptions of lowest rank occurring in every C -justification (instead of all subsumptions occurring in some C -justification).

Definition 5. For $\mathcal{J} \subseteq \mathcal{D}$, let $\mathcal{J}_{\min}^{\mathcal{K}} := \{D \sqsubseteq E \mid r_{\mathcal{K}}(D) \leq r_{\mathcal{K}}(F) \text{ for every } F \sqsubseteq G \in \mathcal{J}\}$. For $C \in \mathcal{L}$, let $\mathcal{J}_{\min}^{\mathcal{K}}(C) := \bigcup_{\mathcal{J} \in \mathcal{J}^{\mathcal{K}}(C)} \mathcal{J}_{\min}^{\mathcal{K}}$.

The intuition can be explained as follows. To make an antecedent C non-exceptional w.r.t. \mathcal{K} , it is necessary to remove at least one element of every C -justification from \mathcal{D} . At the same time, the ranking of subsumptions provides guidance on which subsumptions ought to be removed first (subsumptions with lower ranks are removed first). Combining this, the subsumptions eligible for removal are taken to be precisely those that occur as the lowest ranked subsumptions in some C -justification.

Definition 6. $C \sqsubseteq D$ is in the Minimal Relevant Closure of \mathcal{K} iff it is in the Relevant Closure of \mathcal{K} w.r.t. $\bigcup \mathcal{J}_{\min}^{\mathcal{D}}(C)$.

To see how Minimal Relevant Closure differs from Basic Relevant Closure, we extend Example 1 as follows.

Example 2. In addition to the information in Example 1, we also know that mammalian sickle cells are mammalian red blood cells, that mammalian red blood cells normally have a bioconcave shape, but that mammalian sickle cells normally do not (they normally have a crescent shape). We represent this new information as $\mathcal{T}^2 = \{\text{MSC} \sqsubseteq \text{MRBC}\}$ and $\mathcal{D}^2 = \{\text{MRBC} \sqsubseteq \exists \text{hasS.BC}, \text{MSC} \sqsubseteq \neg \exists \text{hasS.BC}\}$.

To answer the query of whether mammalian sickle cells don't have a nucleus (that is, whether $\text{MSC} \sqsubseteq \neg \exists \text{hasN.T}$) given a KB $\mathcal{K}^2 = \langle \mathcal{T}, \mathcal{D} \rangle$, with $\mathcal{T} = \mathcal{T}^1 \cup \mathcal{T}^2$ and $\mathcal{D} = \mathcal{D}^1 \cup \mathcal{D}^2$, note that there are two MSC-justifications for \mathcal{K} : $\mathcal{J}^1 = \{\text{MRBC} \sqsubseteq \neg \exists \text{hasN.T}, \text{VRBC} \sqsubseteq \exists \text{hasN.T}\}$, and $\mathcal{J}^2 = \{\text{MRBC} \sqsubseteq \exists \text{hasS.BC}, \text{MSC} \sqsubseteq \neg \exists \text{hasS.BC}\}$. Therefore $\text{MSC} \sqsubseteq \neg \exists \text{hasN.T}$ is in the Basic Relevant Closure of \mathcal{K}^2 iff it is in the Relevant Closure of \mathcal{K} w.r.t. R , where R consists of all of \mathcal{D} except for the only irrelevant axiom $\text{VRBC} \sqsubseteq \exists \text{hasCM.T}$ (the only axiom that does not appear in any MSC-justification). It turns out that $\text{MSC} \sqsubseteq \neg \exists \text{hasN.T}$ is not in the Basic Relevant Closure of \mathcal{K}^2 since $\text{MRBC} \sqsubseteq \neg \exists \text{hasN.T}$ is viewed as relevant w.r.t. the query.

To check if $\text{MSC} \sqsubseteq \neg \exists \text{hasN.T}$ is in the Minimal Relevant Closure, note that $\mathcal{J}_{\min}^1 = \{\text{VRBC} \sqsubseteq \exists \text{hasN.T}\}$, and $\mathcal{J}_{\min}^2 = \{\text{MRBC} \sqsubseteq \exists \text{hasS.BC}\}$. Thus, $\text{MSC} \sqsubseteq \neg \exists \text{hasN.T}$ is

in the Minimal Relevant Closure of \mathcal{K} iff it is in the Relevant Closure of \mathcal{K} w.r.t. R , where R consists of everything in \mathcal{D} except for the defeasible subsumptions in the set $\{\text{VRBC} \sqsubseteq \exists \text{hasCM}.\top, \text{MRBC} \sqsubseteq \neg \exists \text{hasN}.\top, \text{MSC} \sqsubseteq \neg \exists \text{hasS}.\text{BC}\}$. And this is the case, since $\text{MRBC} \sqsubseteq \neg \exists \text{hasN}.\top$ is now deemed to be *irrelevant* w.r.t. the query.

To summarise, unlike the case for Rational Closure and Basic Relevant Closure, using Minimal Relevant Closure we can conclude that mammalian sickle cells normally don't have a nucleus. The main reason is that, although Minimal Relevant Closure recognises that mammalian sickle cells are abnormal mammalian red blood cells, the information that mammalian red blood cells do not have a nucleus is deemed to be irrelevant to this abnormality, which means that this defeasible property of mammalian red blood cells are inherited by mammalian sickle cells.

5 Properties of Relevant Closure

The previous sections contain a number of examples showing that both Basic and Minimal Relevant Closure provide better results than Rational Closure. The purpose of this section is to provide a more systematic evaluation. We commence by showing that Minimal Relevant Closure is inferentially stronger than Basic Relevant Closure which, in turn, is inferentially stronger than Rational Closure.

Proposition 2. *If $C \sqsubseteq D$ is in the Rational Closure of a knowledge base \mathcal{K} , then it is in the Basic Relevant Closure of \mathcal{K} (the converse does not always hold). If $C \sqsubseteq D$ is in the Basic Relevant Closure of \mathcal{K} , then it is in the Minimal Relevant Closure of \mathcal{K} (the converse does not always hold).*

It is known that Rational Closure and Preferential Entailment are equivalent w.r.t. the classical subsumptions they contain. The next result shows that this result extends to Basic and Minimal Relevant Closure as well.

Proposition 3. *$C \sqsubseteq D$ is in the Minimal Relevant Closure of a knowledge base \mathcal{K} , iff it is in the Basic Relevant Closure of \mathcal{K} , iff it is in the Rational Closure of \mathcal{K} (iff it is in the Preferential Entailment of \mathcal{K}).*

One of the reasons Proposition 3 is important is that it ensures that Basic and Minimal Relevant Closure are proper generalisations of classical entailment: If $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$ is reduced to classical subsumptions—that is, if \mathcal{K} is well-separated and $\mathcal{D} = \emptyset$ —then Minimal and Basic Relevant Closure coincide with classical entailment.

From a practical point of view, one of the main advantages of both Basic and Minimal Relevant Closure is that, as for Rational Closure, their computation can be reduced to a sequence of classical entailment checks, thereby making it possible to employ existing optimised classical DL reasoners for this purpose. Below we provide high-level algorithms for both versions based on this principle.

The algorithm for Basic Relevant Closure takes as input a well-separated KB $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$, a query $C \sqsubseteq D$, and uses the *partition* procedure which partitions the elements of \mathcal{D} according to ranks. It also assumes the existence of a *justifications* procedure which takes as input a DBox \mathcal{K} , a concept C , and returns the m C -justifications w.r.t. \mathcal{K} . It returns **true** iff the query is in the Basic Relevant Closure of \mathcal{K} .

Algorithm 3: Basic Relevant Closure

Input: A well-separated $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$ and a query $C \sqsubseteq D$
Output: true iff $C \sqsubseteq D$ is in the Basic Relevant Closure of \mathcal{K}

- 1 $(\mathcal{D}_0, \dots, \mathcal{D}_n, n) := \text{partition}(\mathcal{D});$
- 2 $(\mathcal{J}_1, \dots, \mathcal{J}_m, m) := \text{justifications}(\mathcal{K}, C);$
- 3 $\mathcal{J} := \bigcup_{j=1}^m \mathcal{J}_j; i := 0; \mathcal{D}' := \mathcal{D}; X := \emptyset;$
- 4 **while** $X \cap \mathcal{J}_j = \emptyset$ for some $j = 1, \dots, m$ **and** $\mathcal{D}' \neq \emptyset$ **do**
- 5 $\mathcal{D}' := \mathcal{D}' \setminus (\mathcal{J} \cap \mathcal{D}_i);$
- 6 $X := X \cup (\mathcal{J} \cap \mathcal{D}_i); i := i + 1;$
- 7 **return** $\mathcal{T} \models \overline{\mathcal{D}'} \cap C \sqsubseteq D;$

In terms of computational complexity, the big difference between Algorithm 1 and Algorithm 3 is that the latter needs to compute all C -justifications which can involve an exponential number of classical entailment checks [21]. This is in contrast to Algorithm 1 which needs to perform at most a polynomial number of entailment checks. But, since entailment checking for \mathcal{ALC} is EXPTIME-complete, computing the Basic Relevant Closure is EXPTIME-complete as well. From a practical perspective, Horridge [21] has shown that computing justifications is frequently feasible even for large ontologies. We address this issue again in the sections on experimental results and future work.

Next we provide a high-level algorithm for computing Minimal Relevant Closure. Like Algorithm 3, it takes as input a well-separated KB $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$ and a query $C \sqsubseteq D$,

Algorithm 4: Minimal Relevant Closure

Input: A well-separated $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$ and a query $C \sqsubseteq D$
Output: true iff $C \sqsubseteq D$ is in the Minimal Relevant Closure of \mathcal{K}

- 1 $(\mathcal{D}_0, \dots, \mathcal{D}_n, n) := \text{partition}(\mathcal{D});$
- 2 $(\mathcal{J}_1, \dots, \mathcal{J}_m, m) := \text{justifications}(\mathcal{K}, C);$
- 3 **for** $j := 1$ **to** m **do**
- 4 $k := \min(\mathcal{D}_0, \dots, \mathcal{D}_n, \mathcal{J}_j);$
- 5 $\mathcal{M}_j := \mathcal{J}_j \cap \mathcal{D}_k;$
- 6 $\mathcal{M} := \bigcup_{j=1}^m \mathcal{M}_j; i := 0; \mathcal{D}' := \mathcal{D}; X := \emptyset;$
- 7 **while** $X \cap \mathcal{M}_j = \emptyset$ for some $j = 1, \dots, m$ **and** $\mathcal{D}' \neq \emptyset$ **do**
- 8 $\mathcal{D}' := \mathcal{D}' \setminus (\mathcal{M} \cap \mathcal{D}_i);$
- 9 $X := X \cup (\mathcal{M} \cap \mathcal{D}_i); i := i + 1;$
- 10 **return** $\mathcal{T} \models \overline{\mathcal{D}'} \cap C \sqsubseteq D;$

uses the *partition* procedure which partitions the elements of \mathcal{D} according to ranks, and uses the *justifications* procedure which takes as input \mathcal{K} , a concept C , and returns the m C -justifications w.r.t. \mathcal{K} . In addition, it assumes the existence of a *min* procedure which takes as input the partitioned version of \mathcal{D} and any subset of \mathcal{D} , say Y , and returns the smallest j such that $Y \cap \mathcal{D}_j \neq \emptyset$.

Since the only real difference between Algorithm 3 and Algorithm 4 is the use of the *min* procedure, which does not involve any classical entailment check, it follows easily that computing the Minimal Relevant Closure is EXPTIME-complete as well.

To conclude this section we evaluate Basic and Minimal Relevant Closure against the KLM properties of Kraus et al. [23] for rational preferential consequence, translated to DLs.

$$\begin{array}{c}
(\text{Cons}) \top \not\sqsubseteq \perp \quad (\text{Ref}) C \sqsubseteq C \\
(\text{LLE}) \frac{\models C \equiv D, C \sqsubseteq E}{D \sqsubseteq E} \quad (\text{And}) \frac{C \sqsubseteq D, C \sqsubseteq E}{C \sqsubseteq D \sqcap E} \\
(\text{Or}) \frac{C \sqsubseteq E, D \sqsubseteq E}{C \sqcup D \sqsubseteq E} \quad (\text{RW}) \frac{C \sqsubseteq D, \models D \sqsubseteq E}{C \sqsubseteq E} \\
(\text{CM}) \frac{C \sqsubseteq D, C \sqsubseteq E}{C \sqcap D \sqsubseteq E} \quad (\text{RM}) \frac{C \sqsubseteq E, C \not\sqsubseteq \neg D}{C \sqcap D \sqsubseteq E}
\end{array}$$

With the exceptions of Cons, these have been discussed at length in the literature for both the propositional and the DL cases [23, 25, 24, 18] and we shall not do so here. Semantically, Cons corresponds to the requirement that ranked interpretations have non-empty domains. Although these are actually properties of the defeasible subsumption relation \sqsubseteq , they can be viewed as properties of a closure operator as well. That is, we would say that Basic Relevant Closure satisfies the property Ref, for example, whenever $C \sqsubseteq C$ is in the Basic Relevant Closure of \mathcal{D} for every DBox \mathcal{D} and every $C \in \mathcal{L}$.

Proposition 4. *Both Basic Relevant Closure and Minimal Relevant Closure satisfy the properties Cons, Ref, LLE, And, and RW, and do not satisfy Or, CM, and RM.*

While Basic Relevant Closure and Minimal Relevant Closure are inferentially stronger than Rational Closure, and behave well in terms of the examples discussed, their failure to satisfy the formal properties Or, CM and RM is a drawback. We are currently investigating refinements of both Basic Relevant Closure and Minimal Relevant Closure that will satisfy these properties.

6 Experimental Results

In this section we report on preliminary experiments to determine the practical performance of Basic and Minimal Relevant Closure relative to Rational Closure. Our algorithms were implemented and applied to the generated dataset employed by Casini et al. [11]. The DBoxes are binned according to *percentage defeasibility* (ratio of the number of defeasible vs. classical subsumptions) in increments of 10 from 10 to 100, and vary uniformly in size between 150 and 5150 axioms. In addition to the generated DBoxes, we randomly generated a set of DBox queries using terms in their signatures. The task is then to check whether a query is in the Basic (resp. Minimal) Relevant Closure of the DBox and plot its performance relative to Rational Closure. The rankings of each DBox were precomputed because determining the ranking can be viewed as an offline process, and is not the central interest here. Experiments were performed on an Intel Core i7 machine with 4GB of memory allocated to the JVM (Java Virtual Machine). The underlying classical DL reasoning implementation used in our algorithm is Hermit (<http://www.hermit-reasoner.com>). As a preliminary optimisation we prune away

axioms from the rankings that are *irrelevant* to the query according to the notion of *entailment preserving modules* [14]

Results: Overall, the Basic and Minimal Relevant Closure took around one order of magnitude longer to compute than Rational Closure (see Figure 1).

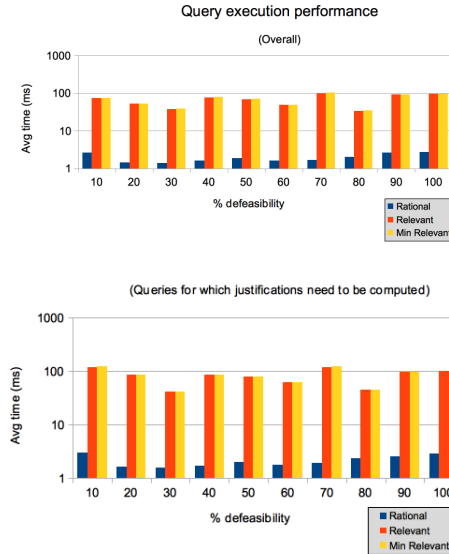


Fig. 1. Average query execution performance of Basic Relevant Closure (in red) vs. Rational Closure (in blue) over the dataset.

The reason for this discrepancy in performance is attributable to the relatively large number of classical entailment checks required to compute the justifications (see Figure 2) for Basic and Minimal Relevant Closure. Rational Closure, on the other hand, does not require to compute justifications and therefore in general is significantly faster. Another contributing factor to this is that HerMiT is not optimised for entailment checks of the form found in Algorithms 3 and 4.

As expected, the performance of Basic (and Minimal) Relevant Closure drastically degrades when it has to compute a large number of justifications. We found that 8% of queries could not be computed in reasonable time. We introduced a timeout of 7000ms, which is one order of magnitude longer than that of the worst case query answering times for Rational Closure (700ms). The timeout accounts, to some extent, for the number of justifications being more or less constant. Despite this, we observe that an average query answering time of 100 milliseconds is promising as an initial result, especially since our algorithms are not highly optimised.

Since the practical feasibility of Basic and Minimal Relevant Closure relies on the justificatory structure of the DBoxes, we plan to investigate the prevalence of justifica-




Fig. 2. Average number of justifications computed for query answering using Basic Relevant Closure.

tions in real-world ontologies. This investigation could reveal the usefulness of these forms of reasoning in such contexts.

Finally, given the minor differences between the algorithms for Basic and Minimal Relevant Closure it is not surprising that the latter behaves very similarly to the former from a performance perspective.

7 Related Work

The semantic underpinnings of our work has its roots in the propositional approach to defeasible reasoning advocated by Lehmann and colleagues [23, 25, 24] and transported to the DL setting by Britz et al. [8, 9] and Giordano et al. [18, 17, 16]. From an algorithmic perspective, Giordano et al. [18] present a tableau calculus for computing Preferential Entailment which relies on KLM-style rules. To our knowledge, this has not been implemented yet. Our work builds on that of Casini and Straccia [12] who describe an algorithm for computing (a slightly different version of) Rational Closure for \mathcal{ALC} , and Britz et al. [7], who refined the Casini-Straccia algorithm to correspond exactly to Rational Closure, and implemented the refined algorithm. Their accompanying experimental results showed that enriching DLs with defeasible subsumption is practically feasible.

Strongly related to our work as well is the approach to defeasible reasoning known as Lexicographic Closure, first proposed by Lehmann [24] for the propositional case, and extended to the DL case by Casini and Straccia [13]. Lukasiewicz [26] also proposed a method that, as a special case, corresponds to a version of Lexicographic Closure. Below we present a description of Lexicographic Closure for $\mathcal{ALC}(\sqsupseteq)$ (space considerations prevent a more detailed description).

Let $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$ be a KB with \mathcal{D} partitioned into $\mathcal{D}_0, \dots, \mathcal{D}_n$. For $\mathcal{D}' \subseteq \mathcal{D}$, let $k_i^{\mathcal{D}'} = |\mathcal{D}_i \cap \mathcal{D}'|$. Let \prec be the lexicographic order on sequences of natural numbers of length $n + 2$. For $\mathcal{D}', \mathcal{D}'' \subseteq \mathcal{D}$, let $\mathcal{D}' \triangleleft \mathcal{D}''$ iff $[k_0^{\mathcal{D}'}, \dots, k_n^{\mathcal{D}'}, k_\infty^{\mathcal{D}'}] \prec [k_0^{\mathcal{D}''}, \dots, k_n^{\mathcal{D}''}, k_\infty^{\mathcal{D}''}]$. For $\mathcal{D}' \subseteq \mathcal{D}$ and $C \in \mathcal{L}$, \mathcal{D}' is a *basis* for C w.r.t. \mathcal{K} iff $\mathcal{T} \not\models \overline{\mathcal{D}'} \sqsupseteq \neg C$ and \mathcal{D}' is maximal w.r.t. the ordering \triangleleft .

Definition 7. For $C, D \in \mathcal{L}$, $C \sqsupseteq D$ is in the Lexicographic Closure of $\mathcal{K} = \langle \mathcal{T}, \mathcal{D} \rangle$ iff for every basis \mathcal{D}' for C w.r.t. \mathcal{K} , $\mathcal{T} \models \bigcap \mathcal{D}' \sqcap C \sqsupseteq D$.

Lexicographic Closure corresponds to what Lehmann [24] refers to *presumptive reasoning* and describes as the reading intended by Reiter’s Default Logic [29]. It satisfies all the KLM properties, and is known to be inferentially stronger than Rational Closure. It turns out to be stronger than Minimal Relevant Closure (and Basic Relevant Closure) as well.

Proposition 5. *If $C \sqsubseteq D$ is in the Minimal Relevant Closure of \mathcal{D} , then it is in the Lexicographic Closure of \mathcal{D} . The converse does not hold.*

Lexicographic Closure is a powerful form of defeasible reasoning and is certainly worth further investigation in the context of DLs. At present, we are not aware of any implementation of Lexicographic Closure, though.

More generally, other proposals for defeasible reasoning include default-style rules in description logics [3, 27], approaches based on circumscription for DLs [6, 5, 4, 30], and approaches that combines an explicit knowledge operator with negation as failure [22, 15]. To our knowledge, the formal properties of the consequence relation of these systems have not been investigated in detail, and none of them have been implemented.

8 Conclusion and Future Work

In this paper we proposed a new approach to defeasible reasoning for DLs based on the *relevance* of subsumptions to a query. We instantiated the approach with two versions of relevance-based defeasible reasoning—Basic Relevant Closure and Minimal Relevant Closure. We showed that both versions overcome some of the limitations of Rational Closure, the best known version of KLM-style defeasible reasoning. We presented experimental results based on an implementation of both Basic Relevant Closure and Minimal Relevant Closure, and compared it with existing results for Rational Closure. The results indicate that both Basic Relevant Closure and Minimal Relevant Closure are only slightly more expensive to compute than Rational Closure.

The relevance-based reasoning proposed in this paper is, to our knowledge, the first attempt to define a form of defeasible reasoning on the use of justifications—a notion on which the area of ontology debugging is based. An obvious extension to the current work is the investigation of relevance-based reasoning other than Basic and Minimal Relevant Closure. We are currently investigating a version that is inferentially stronger than Minimal Relevant Closure.

Here the focus was on defeasible reasoning for DBoxes without reference to ABox assertions. The incorporation of defeasible ABox reasoning into both forms of Relevant Closure presented here is similar to existing approaches for Rational Closure [17, 16, 10] and is left as future work.

Finally, there are two ways to deal with the computational burden associated with Relevant Closure in comparison with Rational Closure. Firstly, there is the option of optimised versions of the current implementations. Secondly, there is the possibility of developing efficient algorithms for approximating Basic or Minimal Relevant Closure, that are guaranteed to be at least as strong as Rational Closure, inferentially speaking. We are pursuing both options.

References

1. F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. Patel-Schneider. *The Description Logic Handbook: Theory, Implementation and Applications*. Cambridge University Press, 2 edition, 2007.
2. F. Baader and B. Hollunder. How to prefer more specific defaults in terminological default logic. In *Proceedings of IJCAI*, pages 669–674. Morgan Kaufmann Publishers, 1993.
3. F. Baader and B. Hollunder. Embedding defaults into terminological knowledge representation formalisms. In *Journal of Automated Reasoning*, pages 306–317. Morgan Kaufmann, 1995.
4. P. Bonatti, M. Faella, and L. Sauro. Defeasible inclusions in low-complexity DLs. *JAIR*, 42:719–764, 2011.
5. P. Bonatti, M. Faella, and L. Sauro. On the complexity of EL with defeasible inclusions. In *Proceedings of IJCAI*, pages 762–767, 2011.
6. P. Bonatti, C. Lutz, and F. Wolter. The complexity of circumscription in description logic. *JAIR*, 35:717–773, 2009.
7. K. Britz, G. Casini, T. Meyer, K. Moodley, and I. Varzinczak. Ordered interpretations and entailment for defeasible description logics. Technical report, CAIR, CSIR Meraka and UKZN, South Africa, 2013.
8. K. Britz, J. Heidema, and T. Meyer. Semantic preferential subsumption. In J. Lang and G. Brewka, editors, *Proceedings of KR*, pages 476–484. AAAI Press/MIT Press, 2008.
9. K. Britz, T. Meyer, and I. Varzinczak. Semantic foundation for preferential description logics. In D. Wang and M. Reynolds, editors, *Proceedings of AI*, number 7106 in LNAI, pages 491–500. Springer, 2011.
10. G. Casini, T. Meyer, K. Moodley, and I. Varzinczak. Nonmonotonic reasoning in description logics: Rational closure for the Abox. In *Proceedings of DL*, pages 600–615, 2013.
11. G. Casini, T. Meyer, K. Moodley, and I. Varzinczak. Towards practical defeasible reasoning for description logics. In *Proceedings of DL*, pages 587–599, 2013.
12. G. Casini and U. Straccia. Rational closure for defeasible description logics. In T. Janhunen and I. Niemelä, editors, *Proceedings of JELIA*, number 6341 in LNCS, pages 77–90. Springer, 2010.
13. G. Casini and U. Straccia. Lexicographic closure for defeasible description logics. In *Proc of Australasian Ontology Workshop*, 2012.
14. B. Cuenca Grau, I. Horrocks, Y. Kazakov, and U. Sattler. Modular reuse of ontologies: Theory and practice. *JAIR*, 31:273–318, 2008.
15. F. M. Donini, D. Nardi, and R. Rosati. Description logics of minimal knowledge and negation as failure. *TOCL*, 3(2):177–225, 2002.
16. L. Giordano, V. Gliozzi, and N. Olivetti. Minimal model semantics and rational closure in description logics. *Proceedings of DL*, 2013.
17. L. Giordano, V. Gliozzi, N. Olivetti, and G. L. Pozzato. Rational Closure in Description Logics of Typicality. *IAF*, 2013.
18. L. Giordano, N. Olivetti, V. Gliozzi, and G.L. Pozzato. $\mathcal{ALC} + T$: a preferential extension of description logics. *Fundamenta Informaticae*, 96(3):341–372, 2009.
19. L. Giordano, N. Olivetti, V. Gliozzi, and G.L. Pozzato. A minimal model semantics for non-monotonic reasoning. In L. Fariñas del Cerro, A. Herzig, and J. Mengin, editors, *Proceedings of JELIA*, number 7519 in LNCS, pages 228–241. Springer, 2012.
20. S.O. Hansson. *A Textbook of Belief Dynamics: Theory Change and Database Updating*. Kluwer Academic Publishers, 1999.
21. M. Horridge. *Justification based explanation in ontologies*. The University of Manchester, 2011.

22. P. Ke and U. Sattler. Next Steps for Description Logics of Minimal Knowledge and Negation as Failure. In *Proceedings of DL*, 2008.
23. S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44:167–207, 1990.
24. D. Lehmann. Another perspective on default reasoning. *Annals of Mathematics and Artificial Intelligence*, 15(1):61–82, 1995.
25. D. Lehmann and M. Magidor. What does a conditional knowledge base entail? *Artificial Intelligence*, 55:1–60, 1992.
26. T. Lukasiewicz. Expressive probabilistic description logics. *Artificial Intelligence*, 172(6-7):852–883, 2008.
27. L. Padgham and T. Zhang. A terminological logic with defaults: A definition and an application. In *Proceedings of IJCAI*, pages 662–668. Morgan Kaufmann, 1994.
28. J. Quantz. A preference semantics for defaults in terminological logics. In *Proceedings of KR*, pages 294–305, 1992.
29. R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13(1-2):81–132, April 1980.
30. K. Sengupta, A. Alfa Krisnadhi, and P. Hitzler. Local closed world semantics: Grounded circumscription for OWL. In *Proceedings of ISWC*, pages 617–632, 2011.