Concept Model Semantics for DL Preferential Reasoning

Katarina Britz[1,2], Thomas Meyer[1,2], and Ivan Varzinczak[1,2]

[1] CSIR Meraka Institute, Pretoria, South Africa
[2] University of KwaZulu-Natal, Durban, South Africa
{arina.britz,tommie.meyer,ivan.varzinczak}@meraka.org.za

Abstract. The preferential and rational consequence relations first studied by Lehmann and colleagues play a central role in non-monotonic reasoning, not least because they provide the foundation for the determination of the important notion of rational closure. Although they can be applied directly to a large variety of logics, these constructions suffer from the limitation that they are largely propositional in nature. One of the main obstacles in moving beyond the propositional case has been the lack of a formal semantics which appropriately generalizes the preferential and ranked models of Lehmann et al. In this paper we propose a semantics to fill that gap for description logics, an important class of decidable fragments of first-order logic. Our semantics replaces the propositional valuations used in the models of Lehmann et al. with structures we refer to as *concept models*. We prove representation results for the description logic $\mathcal{ALC}$ for both preferential and rational consequence relations. We argue that our semantics paves the way for extending preferential and rational consequence, and therefore also rational closure, to a whole class of logics that have a semantics defined in terms of first-order relational structures.

1 Introduction

There has by now been quite a substantial number of attempts to incorporate defeasible reasoning in logics other than propositional logic. One such endeavor, and the broad focus of this paper, has been to extend the influential version of preferential reasoning first studied by Lehmann et al. [7, 9] to logics beyond the propositional. A stumbling block to this end has been that research on preferential reasoning has really only reached maturity in a propositional context, whereas many logics of interest have more structure. A generally accepted semantics for first-order preferential reasoning, with corresponding syntactic proof system or characterization, does not yet exist. The first tentative exploration of preferential predicate logics by Lehmann et al. didn't fly (pun intended), primarily because propositional logic was sufficiently expressive for the non-monotonic reasoning community at the time, and first-order logic introduced too much complexity [8]. But this changed with the surge of interest in description logics as knowledge representation formalism. Description logics (DLs) [1] are decidable fragments of first-order logic, and are ideal candidates for the kind of extension to preferential reasoning we have in mind: the notion of subsumption present in all DLs is a natural candidate for defeasibility, while at the same time, the restricted expressivity of DLs ensures that attempts to introduce preferential reasoning are not hampered by the complexity of full first-order logic. The aim of this paper is to extend the work of Lehmann et al. [7, 9] beyond propositional logic without moving to full first-order logic. We restrict our attention to the description logic $\mathcal{ALC}$ here, but the results are broadly applicable to other DLs, as well as other similarly structured logics such as logics of action and logics of knowledge and belief.

The central question answered in this paper is how the existing semantics for both preferential and rational propositional non-monotonic consequence relations should be generalized to languages with more structure, and in particular, to $\mathcal{ALC}$. Specifically, what is the meaning of a preferential (or rational) subsumption statement $C \sqsubseteq\!\!\!\!\sim D$ — what properties should it have, and what is its corresponding formal semantics? The main results of this paper, and our answers to these questions, are the two representation results presented in Theorems 3 and 4, respectively. Key to the establishment of these results are the notions of a *concept model*, which gives a reading of the meaning of concepts suitable for our purposes, and a *DL preferential model*, giving meaning to non-monotonic subsumption statements. The latter generalizes the notion of a propositional preferential model in terms of concept models.

The rest of the paper is structured as follows. In Section 2 we give a brief account of the work on preferential and rational subsumption for the propositional case as developed by Lehmann and colleagues. Section 3 is the heart of the paper in which we define the semantics for both preferential and rational subsumption for $\mathcal{ALC}$ and prove representation results for both. Importantly, the representation results provided here are with respect to the corresponding *propositional* properties. From this we conclude that the semantics we present here forms the foundation of a semantics for preferential and rational consequence for a whole class of DLs and related logics and provides a natural and intuitive semantic framework on which to base such work. In Section 4 we use the fundamental results of the previous section to show that the notions of propositional preferential entailment and rational closure can be 'lifted' to the case for DLs, specifically $\mathcal{ALC}$. In Section 5 we discuss related results. We conclude with Section 6 in which we also discuss future work.

We assume that the reader is familiar with description logics. For more details on description logics in general, and the description logic $\mathcal{ALC}$ in particular, the reader is referred to the DL handbook [1].

2 Propositional Preferential Consequence

In this section we give a brief introduction to propositional preferential and rational consequence, as initially defined by Kraus et al. [7]. A propositional defeasible consequence relation $\mid\!\sim$ is defined as a binary relation on formulas

$\alpha, \beta, \gamma, \ldots$ of an underlying (possibly infinitely generated) propositional logic equipped with a standard propositional entailment relation $\models$ [7]. $\mathrel{|\!\sim}$ is said to be *preferential* if it satisfies the following set of properties:

$$\text{(Ref)}\ \alpha \mathrel{|\!\sim} \alpha \qquad \text{(LLE)}\ \frac{\alpha \equiv \beta, \quad \alpha \mathrel{|\!\sim} \gamma}{\beta \mathrel{|\!\sim} \gamma} \qquad \text{(And)}\ \frac{\alpha \mathrel{|\!\sim} \beta, \quad \alpha \mathrel{|\!\sim} \gamma}{\alpha \mathrel{|\!\sim} \beta \wedge \gamma}$$

$$\text{(RW)}\ \frac{\alpha \mathrel{|\!\sim} \beta, \quad \beta \models \gamma}{\alpha \mathrel{|\!\sim} \gamma} \qquad \text{(Or)}\ \frac{\alpha \mathrel{|\!\sim} \gamma, \quad \beta \mathrel{|\!\sim} \gamma}{\alpha \vee \beta \mathrel{|\!\sim} \gamma} \qquad \text{(CM)}\ \frac{\alpha \mathrel{|\!\sim} \beta, \quad \alpha \mathrel{|\!\sim} \gamma}{\alpha \wedge \beta \mathrel{|\!\sim} \gamma}$$

The semantics of (propositional) preferential consequence relations is in terms of *preferential models*; these are partially ordered structures with states labeled by propositional valuations. We shall make this terminology more precise in Section 3, but it essentially allows for a partial order on states, with states lower down in the order being more preferred than those higher up. Given a preferential model $\mathscr{P}$, a pair $\alpha \mathrel{|\!\sim} \beta$ is in the consequence relation defined by $\mathscr{P}$ iff the minimal states (according to the partial order) of all those states labeled by valuations that are propositional models of α, are also labeled by propositional models of β. The representation theorem for preferential consequence relations then states:

Theorem 1 (Kraus et al. [7]). *A defeasible consequence relation is a preferential consequence relation iff it is defined by some preferential model.*

If, in addition to the properties of preferential consequence, $\mathrel{|\!\sim}$ also satisfies the following Rational Monotony property, it is said to be a *rational* consequence relation:

$$\text{(RM)}\ \frac{\alpha \mathrel{|\!\sim} \beta, \quad \alpha \not\mathrel{|\!\sim} \neg\gamma}{\alpha \wedge \gamma \mathrel{|\!\sim} \beta}$$

The semantics of rational consequence relations is in terms of *ranked* preferential models, i.e., preferential models in which the preference order is *modular*:

Definition 1. *Given a set S, $\prec \subseteq S \times S$ is modular iff $\prec$ is a partial order on S, and there is a ranking function $rk : S \mapsto \mathbb{N}$ s.t. for every $s, s' \in S$, $s \prec s'$ iff $rk(s) < rk(s')$.*

The representation theorem for rational consequence relations then states:

Theorem 2 (Lehmann and Magidor [9]). *A defeasible consequence relation is a rational consequence relation iff it is defined by some ranked model.*

3 Semantics for DL Preferential Consequence

It has been argued elsewhere that description logics are ideal candidates for the extension of propositional preferential consequence since the notion of subsumption in DLs lends itself naturally to defeasibility [3, 6, 4]. The basic idea is to reinterpret defeasible consequence of the form $\alpha \mathrel{|\!\sim} \beta$ as *defeasible subsumption* of the form $C \precsim D$, where C and D are DL concepts, and classical entailment $\models$ as DL subsumption $\sqsubseteq$. The properties of preferential consequence from Section 2 are then immediately applicable.

Definition 2. *A subsumption relation $\precsim \subseteq \mathcal{L} \times \mathcal{L}$ is a preferential subsumption relation iff it satisfies the properties (Ref), (LLE), (And), (RW), (Or), and (CM), with propositional entailment replaced by classical DL subsumption. $\precsim$ is a rational subsumption relation iff in addition to being a preferential subsumption relation, it also satisfies the property (RM).*

However, up until now it has not been clear how to best generalize the propositional semantics for the DL case. Since DLs have a standard first-order semantics, the obvious generalization from a technical perspective is to replace the propositional valuations in preferential models with first-order interpretations. Intuitively, this also turns out to be a natural generalization of the propositional setting, with the notion of normal first-order interpretation characterizing a given concept replacing the propositional notion of normal worlds satisfying a given proposition. Formally, our semantics is based on the notion of a *concept model*, which is analogous to that of a Kripke model in modal logic [2]:

Definition 3 (Concept Model). *A concept model is a tuple $\mathscr{M} = \langle W, R, V \rangle$ where W is a set of possible worlds, $R = \langle R_1, \ldots, R_n \rangle$, where each $R_i \subseteq W \times W$, $1 \leq i \leq |\mathsf{N}_{\mathscr{R}}|$, and $V \colon W \mapsto 2^{\mathsf{N}_{\mathscr{C}}}$ is a valuation function.*

Observe that the valuation function V can be viewed as a propositional valuation with propositional atoms replaced by concept names. From the definition of satisfaction in a concept model below it is then clear that, within the context of a concept model, a world occurring in that concept model is a proper generalization of a propositional valuation.

Definition 4 (Satisfaction). *Given $\mathscr{M} = \langle W, R, V \rangle$ and $w \in W$:*

- $\mathscr{M}, w \Vdash \top$;
- $\mathscr{M}, w \Vdash A$ *iff* $A \in V(w)$;
- $\mathscr{M}, w \Vdash C \sqcap D$ *iff* $\mathscr{M}, w \Vdash C$ *and* $\mathscr{M}, w \Vdash D$;
- $\mathscr{M}, w \Vdash \neg C$ *iff* $\mathscr{M}, w \nVdash C$;
- $\mathscr{M}, w \Vdash \exists r_i.C$ *iff there is* $w' \in W$ *s.t.* $(w, w') \in R_i$ *and* $\mathscr{M}, w' \Vdash C$.

Let $\mathscr{U}$ denote the set of all pairs $(\mathscr{M}, w)$ where $\mathscr{M} = \langle W, R, V \rangle$ is a concept model and $w \in W$.

Worlds are, loosely speaking, interpreted DL objects. And while this correspondence holds technically (from the correspondence between $\mathcal{ALC}$ and multimodal logic K [13]), a possible worlds reading of the meaning of a concept is also more intuitive in the current context, since this leads to a preference order on rich first-order structures, rather than on interpreted objects. This is made precise below.

Let S be a set, the elements of which are called *states*. Let $\ell : S \mapsto \mathscr{U}$ be a *labeling function* mapping every state to a pair $(\mathscr{M}, w)$ where $\mathscr{M} = \langle W, R, V \rangle$ is a concept model s.t. $w \in W$. Let $\prec$ be a binary relation on S. Given $C \in \mathcal{L}$, we say that $s \in S$ *satisfies* C (written $s \models\!\equiv C$) iff $\ell(s) \Vdash C$, i.e., $\mathscr{M}, w \Vdash C$. We define $\widehat{C} = \{s \in S \mid s \models\!\equiv C\}$. $\widehat{C}$ is *smooth* iff each $s \in \widehat{C}$ is either $\prec$-minimal in $\widehat{C}$, or there is $s' \in \widehat{C}$ s.t. $s' \prec s$ and s' is $\prec$-minimal in $\widehat{C}$. We say that S satisfies the smoothness condition iff for every $C \in \mathcal{L}$, $\widehat{C}$ is smooth.

We are now ready for our definition of preferential model.

Definition 5 (Preferential Model). *A preferential model is a triple $\mathscr{P} = \langle S, \ell, \prec \rangle$ where S is a set of states satisfying the smoothness condition, ℓ is a labeling function mapping states to elements of $\mathscr{U}$, and $\prec$ is a strict partial order on S, i.e., $\prec$ is irreflexive and transitive.*

These formal constructions closely resemble those of Kraus et al. [7] and of Lehmann and Magidor [9], the difference being that propositional valuations are replaced with elements of the set $\mathscr{U}$.

Definition 6 (Preferential Subsumption). *Given concepts $C, D \in \mathcal{L}$ and a preferential model $\mathscr{P} = \langle S, \ell, \prec \rangle$, we say that C is preferentially subsumed by D in $\mathscr{P}$ (denoted $C \precsim_{\mathscr{P}} D$) iff every $\prec$-minimal state $s \in \widehat{C}$ is s.t. $s \in \widehat{D}$.*

We are now in a position to prove one of the central results of this paper.

Theorem 3. *A defeasible subsumption relation is a preferential subsumption relation iff it is defined by some preferential model.*

The significance of this is that the representation result is proved with respect to the same set of properties used to characterize propositional preferential consequence. We therefore argue that preferential models, as we have defined them, provide the foundation for a semantics for preferential (and rational) subsumption for a whole class of DLs and related logics. We do not claim that this is *the* appropriate notion of preferential subsumption for $\mathcal{ALC}$, but rather that it describes the basic framework within which to investigate such a notion.

In order to obtain a similar result for rational subsumption, we restrict ourselves to those preferential models in which $\prec$ is a modular order on states (cf. Definition 1):

Definition 7 (Ranked Model). *A ranked model $\mathscr{P}_r$ is a preferential model $\langle S, \ell, \prec \rangle$ in which $\prec$ is modular.*

Since ranked models are preferential models, the notion of rational subsumption is as in Definition 6. We can then state the following result:

Theorem 4. *A defeasible subsumption relation is a rational subsumption relation iff it is defined by some ranked model.*

4 Rational Closure

One of the primary reasons for defining non-monotonic consequence relations of the kind we have presented above is to get at a notion of *defeasible entailment*: Given a set of subsumption statements of the form $C \precsim D$ or $C \sqsubseteq D$, which other subsumption statements, defeasible and classical, should one be able to derive from this? It can be shown that classical subsumption statements of the form $C \sqsubseteq D$ can be encoded as defeasible subsumption statements of the form $C \sqcap \neg D \precsim \bot$. For the remainder of this paper we shall therefore concern ourselves only with *finite* sets of defeasible subsumption statements, and refer to these as *defeasible TBoxes*, denoted $\mathcal{T}$. We permit ourselves the freedom to include classical subsumption statements of the form $C \sqsubseteq D$ in a defeasible TBox, with the understanding that it is an encoding of the defeasible subsumption statement $C \sqcap \neg D \precsim \bot$.

Our aim in this section is to show that the results for the propositional case [9] with respect to the question above can be 'lifted' to $\mathcal{ALC}$. We provide here appropriate notions of preferential entailment and rational closure. It must be emphasized that the results obtained in this section rely heavily on similar results obtained by Lehmann and Magidor [9] for the propositional case, and the semantics for preferential and rational subsumption presented in Section 3. Similar to the results of that section, our claim is not that the versions of preferential and rational closure here are *the* appropriate ones for $\mathcal{ALC}$. In fact, our conjecture is that they are *not*, due to their propositional nature. However, we claim that they provide the appropriate springboard from which to investigate more appropriate versions, for $\mathcal{ALC}$, as well as for other DLs and related logics.

The version of rational closure defined here provides us with a strict generalization of classical entailment for $\mathcal{ALC}$ TBoxes in which the expressivity of $\mathcal{ALC}$ is enriched with the ability to make defeasible subsumption statements. For example, consider the defeasible $\mathcal{ALC}$ TBox:

$$\mathcal{T} = \{BM \sqsubseteq M, VM \sqsubseteq M, M \precsim \neg F, BM \precsim F\},$$

where BM abbreviates the concept *BacterialMeningitis*, M stands for *Meningitis*, VM for *viralMeningitis*, and F abbreviates *FatalDisease*. One should be able to conclude that viral meningitis is usually non-fatal ($VM \precsim \neg F$). On the other hand, we should not conclude that fatal versions of meningitis are usually bacterial ($F \sqcap M \precsim BM$), nor, for that matter, that fatal versions of meningitis are usually *not* bacterial ones ($F \sqcap M \precsim \neg BM$).

Armed with the notion of a preferential model (cf. Section 3) we define preferential entailment for $\mathcal{ALC}$ as follows.

Definition 8. *$C \precsim D$ is preferentially entailed by a defeasible TBox $\mathcal{T}$ iff for every preferential model $\mathscr{P}$ in which $E \precsim_{\mathscr{P}} F$ for every $E \precsim F \in \mathcal{T}$, it is also the case that $C \precsim_{\mathscr{P}} D$.*

Firstly, we can show that preferential entailment is well-behaved and coincides with *preferential closure* under the properties of preferential subsumption (i.e., the intersection of all preferential subsumption relations containing a defeasible TBox). More precisely, let $\mathcal{T}$ be a defeasible TBox. Then the set of defeasible subsumption statements preferentially entailed by $\mathcal{T}$, viewed as a binary relation on concepts, is a preferential subsumption relation. Furthermore, a defeasible subsumption statement is preferentially entailed by $\mathcal{T}$ iff it is in the preferential closure of $\mathcal{T}$.

From this it follows that if we use preferential entailment, the meningitis example can be formalized by letting $\mathcal{T} = \{BM \sqsubseteq M\ VM \sqsubseteq M,\ M \precsim \neg F,\ BM \precsim \neg F\}$. However, $VM \precsim \neg F$ is *not* preferentially entailed by $\mathcal{T}$ above (we cannot conclude that viral meningitis is usually not fatal) and preferential entailment is thus generally too weak. We therefore move to rational subsumption relations.

The first attempt to do so is to use a definition similar to that employed for preferential entailment: $C \precsim D$ is rationally entailed by a defeasible TBox $\mathcal{T}$ iff for every ranked model $\mathscr{P}_r$ in which $E \precsim_{\mathscr{P}_r} F$ for every $E \precsim F \in \mathcal{T}$, it is also the case that $C \precsim_{\mathscr{P}_r} D$. However, this turns out to be *exactly* equivalent to preferential entailment. Therefore, if the set of defeasible subsumption statements obtained as such is viewed as a binary relation on concepts, the result is a preferential subsumption relation and is not, in general, a rational consequence relation.

The above attempt to define rational entailment is thus not acceptable. Instead, in order to arrive at an appropriate notion of (rational) entailment we first define a preference ordering on rational subsumption relations, with relations further down in the ordering interpreted as more preferred.

Definition 9. *Let $\precsim_0$ and $\precsim_1$ be rational subsumption relations. $\precsim_0$ is preferable to $\precsim_1$ (written $\precsim_0 \ll \precsim_1$) iff*

- *there is $C \precsim D \in \precsim_1 \setminus \precsim_0$ s.t. for all E s.t. $E \sqcup C \precsim_0 \neg C$ and for all F s.t. $E \precsim_0 F$, we also have $E \precsim_1 F$; and*
- *for every $E, F \in \mathcal{L}$, if $E \precsim F$ is in $\precsim_0 \setminus \precsim_1$, then there is an assertion $G \precsim H$ in $\precsim_1 \setminus \precsim_0$ s.t. $G \sqcup H \precsim_1 \neg H$.*

Space considerations prevent us from giving a detailed motivation for $\ll$ here, but it is essentially the motivation for the same ordering for the propositional case provided by Lehmann and Magidor [9]. Given a defeasible TBox $\mathcal{T}$, the idea is now to define rational entailment as the most preferred (w.r.t. $\ll$) of all those rational subsumption relations which include $\mathcal{T}$.

Lemma 1. *Let $\mathcal{T}$ be a (finite) defeasible TBox and let $\mathcal{R}$ be the class of all rational subsumption relations which include $\mathcal{T}$. There is a unique rational subsumption relation in $\mathcal{R}$ which is preferable to all other elements of $\mathcal{R}$ w.r.t. $\ll$.*

This puts us in a position to define an appropriate form of (rational) entailment for defeasible TBoxes:

Definition 10. *Let $\mathcal{T}$ be a defeasible TBox. The rational closure of $\mathcal{T}$ is the (unique) rational subsumption relation which includes $\mathcal{T}$ and is preferable (w.r.t. $\ll$) to all other rational subsumption relations including $\mathcal{T}$.*

It can be shown that $VM \precsim \neg F$ is in the rational closure of $\mathcal{T}$ (we can conclude viral meningitis is usually not fatal), but that neither $F \sqcap M \precsim BM$ nor $F \sqcap M \precsim \neg BM$ is.

We conclude this section with a result which can be used to define an algorithm for computing the rational closure of a defeasible TBox $\mathcal{T}$. For this we first need to define a ranking of concepts w.r.t. $\mathcal{T}$ which, in turn, is based on a notion of exceptionality. A concept C is said to be *exceptional* for a defeasible TBox $\mathcal{T}$ iff $\mathcal{T}$ *preferentially* entails $\top \precsim \neg C$. A defeasible subsumption statement $C \precsim D$ is exceptional for $\mathcal{T}$ if and only if its antecedent C is exceptional for $\mathcal{T}$.

It turns out that checking for exceptionality can be reduced to classical subsumption checking.

Lemma 2. *Given a defeasible TBox $\mathcal{T}$, let $\mathcal{T}^{\sqsubseteq}$ be its classical counterpart in which every defeasible subsumption statement of the form $D \precsim E$ in $\mathcal{T}$ is replaced by $D \sqsubseteq E$. C is exceptional for $\mathcal{T}$ iff $\top \sqsubseteq \neg C$ is classically entailed by $\mathcal{T}^{\sqsubseteq}$.*

Let $E(\mathcal{T})$ denote the subset of $\mathcal{T}$ containing statements that are exceptional for $\mathcal{T}$. We define a non-increasing sequence of subsets of $\mathcal{T}$ as follows: $\mathcal{E}_0 = \mathcal{T}$, and for $i > 0$, $\mathcal{E}_i = E(\mathcal{E}_{i-1})$. Clearly there is a smallest integer k s.t. for all $j \geq k$, $\mathcal{E}_j = \mathcal{E}_{j+1}$. From this we define the *rank* of a concept w.r.t. $\mathcal{T}$: $r_{\mathcal{T}}(C) = k - i$, where i is the smallest integer s.t. C is not exceptional for $\mathcal{E}_i$. If C is exceptional for $\mathcal{E}_k$ (and therefore exceptional for all $\mathcal{E}$s), then $r_{\mathcal{T}}(C) = 0$. Intuitively, the lower the rank of a concept, the more exceptional it is w.r.t. the TBox $\mathcal{T}$.

Theorem 5. *Let $\mathcal{T}$ be a defeasible TBox. The* rational closure *of $\mathcal{T}$ is the set of defeasible subsumption statements $C \precsim D$ s.t. either $r_{\mathcal{T}}(C) > r_{\mathcal{T}}(C \sqcap \neg D)$, or $r_{\mathcal{T}}(C) = 0$ (in which case $r_{\mathcal{T}}(C \sqcap \neg D) = 0$ as well).*

From this result it is easy to construct a (naïve) decidable algorithm to determine whether a given defeasible subsumption statement is in the rational closure of a defeasible TBox $\mathcal{T}$. Also, if checking for exceptionality is assumed to take constant time, the algorithm is quadratic in the size of $\mathcal{T}$. Given that exceptionality reduces to subsumption checking in $\mathcal{ALC}$ which is ExpTime-complete, it immediately follows that checking whether a given defeasible subsumption statement is in the rational closure of $\mathcal{T}$ is an ExpTime-complete problem. This result is closely related to a result by Casini et al. [4] which we refer to again in the next section.

5 Related Work

Quantz and Ryan [11, 12] were probably the first to consider the lifting of non-monotonic reasoning formalisms to a DL setting. They propose a general framework for Preferential Default Description Logics (PDDL) based on an $\mathcal{ALC}$-like language by introducing a version of default subsumption and proposing a semantics for it. Their semantics is based on a simplified version of standard DL interpretations in which all domains are assumed to be finite and the unique name assumption holds for object names. Their framework is thus much more restrictive than ours. They focus on a version of entailment which they refer to as preferential entailment, but which is to be distinguished from the version of preferential entailment we have presented in this paper. We shall refer to their version as *Q-preferential entailment*.

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

Britz et al. [3] and Giordano et al. [6] use typicality orderings on *objects* in first-order domains to define versions of defeasible subsumption for $\mathcal{ALC}$ and extensions thereof. Both approaches propose specific non-monotonic consequence relations, and hence their semantic constructions are special cases of the more general framework we have provided here. In contrast, we provide a general semantic framework which is relevant to all logics with a possible worlds semantics. This is because our preference semantics is not defined in terms of orders on interpreted DL objects relative to given concepts, but rather in terms of a single order on relational structures. Our semantics for defeasible subsumption yields a single order at the meta level, rather than ad hoc relativized orders at the object level.

Casini and Straccia [4] recently proposed a syntactic operational characterization of rational closure in the context of description logics, based on classical entailment tests only, and thus amenable to implementation. Their work is based on that of Lehmann and Magidor [9], Freund [5] and Poole [10], and represents an important building block in the extension of preferential consequence to description logics. However, this work lacks a semantics, and we can only at present conjecture that the rational closure produced by their algorithm coincides with the notion of the rational closure of a defeasible TBox presented in this paper.

6 Conclusion and Future Work

The main contribution of this paper is the provision of a natural and intuitive formal semantics for preferential and rational subsumption for the description logic $\mathcal{ALC}$. We claim that our semantics provides the foundation for extending preferential reasoning in at least three ways. Firstly, as we have seen in Section 4, it allows for the 'lifting' of preferential entailment and rational closure from the propositional case to the case for $\mathcal{ALC}$. Without the semantics such a lifting may be possible in principle, but will be very hard to prove formally. Secondly, it paves the way for defining similar results for other DLs, as well as other similarly structured logics, such as logics of action and belief. We are at present investigating similar notions for logics of action. And thirdly, it provides the tools to tighten up the versions of preferential and rational subsumption for $\mathcal{ALC}$ presented in this paper in order to truly move beyond the propositional. The latter point is the obvious one to pursue first when it comes to future work. Below we provide some initial ideas on moving beyond propositional properties. The value added by the semantics is the ability it provides to test whether appropriate constraints on the orderings in ranked models can be found that matches the new properties.

Consider the following defeasible TBox, which is a slightly modified version of our previous meningitis example:

$$\mathcal{T} = \{BM \sqsubseteq M, M \precsim \neg F, BM \precsim F, \top \sqsubseteq \forall cm.F\},$$

where BM, M, and F are as before, and cm abbreviates the role *causaMortis*. The last statement encodes the classical subsumption statement that all causes of death are fatal.

It is easily verified that $\exists pc.BM \sqsubseteq \exists pc.M$ is in the rational closure of $\mathcal{T}$ (where pc abbreviates the role *potentiallyCauses*), and so it should be since it is entailed by $BM \sqsubseteq M$. We would also expect to conclude $\exists pc.M \precsim \exists pc.\neg F$ from $\mathcal{T}$ since it contains $M \precsim \neg F$. However, there is no propositional property to guarantee the latter. This prompts us to consider the following property:

$$\frac{C \precsim D}{\exists r.C \precsim \exists r.D} \text{ (Naïve Role Introduction)}$$

Arguably, then, if meningitis is usually non-fatal, then potential causes of meningitis are usually potential causes of something non-fatal.

But there are problems with this reasoning. The following example makes this explicit: From $M \precsim \neg F$, Naïve Role Introduction also allows us to conclude that $\exists cm.M \precsim \exists cm.\neg F$. So usually, fatal cases of meningitis are fatal cases of something non-fatal. This is clearly counter-intuitive. Intuitively, cm usually relates to an abnormal type of meningitis, such as bacterial meningitis, which is usually fatal. An additional blocking mechanism is therefore needed to prevent the rule from being applied when the entire range of the role r is abnormal with respect to C. In order to provide such a mechanism, we need to go beyond $\mathcal{ALC}$, and include the ability to express *role inverses*.[3] We then have the following Role Introduction property:

$$\text{(RI)} \ \frac{C \precsim D, \quad C \not\precsim \forall r^{-}.\bot}{\exists r.C \precsim \exists r.D}$$

The effect of the premise $C \not\precsim \forall r^{-}.\bot$ is to block application of the rule if C is normally disjoint from the range of r. On the other hand, if normally C overlaps with the range of r, it follows that $\exists r.C \precsim \exists r.D$.

[3] Given a role name r, the role inverse of r is denoted by r^{-}. For an interpretation $\mathcal{I}$, $(r^{-})^{\mathcal{I}} = \{(y, x) \mid (x, y) \in r^{\mathcal{I}}\}$.

Now consider again the example above. The intention of role cm is modeled by $\top \sqsubseteq \forall cm.F$, the intuition being that only something fatal can be the cause of death. It then follows classically that $\neg F \sqsubseteq \forall cm^{-}.\bot$, and by (RW) that $M \precsim \forall cm^{-}.\bot$. This property (RI) is therefore blocked for the statement $M \precsim \neg F$. Note that (RI) applied to pc is not blocked, as we do not have that $M \precsim \forall pc^{-}.\bot$. (RI) applied to cm is also not blocked for $BM \precsim F$. The interesting thing about (RI) is that it does *not* hold for the preferential closure of a TBox, whereas it *does* hold for the rational closure.

This illustrates that there are intuitively appealing properties characterizing rational DL-entailment that merit further investigation.

References

1. Baader, F., Calvanese, D., McGuinness, D., Nardi, D., Patel-Schneider, P. (eds.): The Description Logic Handbook. Cambridge, 2 edn. (2007)
2. Blackburn, P., van Benthem, J., Wolter, F.: Handbook of Modal Logic. Elsevier (2006)
3. Britz, K., Heidema, J., Meyer, T.: Semantic preferential subsumption. In: Proc. KR. pp. 476–484. AAAI Press/MIT Press (2008)
4. Casini, G., Straccia, U.: Rational closure for defeasible description logics. In: Proc. of JELIA. pp. 77–90. Springer-Verlag (2010)
5. Freund, M.: Preferential reasoning in the perspective of Poole default logic. Artificial Intelligence 98, 209–235 (1998)
6. Giordano, L., Olivetti, N., Gliozzi, V., Pozzato, G.: $\mathcal{ALC}+T$: a preferential extension of description logics. Fund. Informatica 96(3), 341–372 (2009)
7. Kraus, S., Lehmann, D., Magidor, M.: Nonmonotonic reasoning, preferential models and cumulative logics. Artificial Intelligence 44, 167–207 (1990)
8. Lehmann, D., Magidor, M.: Preferential logics: the predicate calculus case. In: Proc. of TARK. pp. 57–72 (1990)
9. Lehmann, D., Magidor, M.: What does a conditional knowledge base entail? Artificial Intelligence 55, 1–60 (1992)
10. Poole, D.: A logical framework for default reasoning. Artificial Intelligence 36, 27–47 (1988)
11. Quantz, J.J.: A Preference Semantics for Defaults in Terminological Logics. In: Proc. of KR. pp. 294–305 (1992)
12. Quantz, J., Ryan, M.: Preferential Default Description Logics. Tech. rep., TU Berlin (1993), www.tu-berlin.de/fileadmin/fg53/KIT-Reports/r110.pdf
13. Schild, K.: A correspondence theory for terminological logics: Preliminary report. In: Proc. of IJCAI. pp. 466–471 (1991)