

Preferential Role Restrictions

Katarina Britz, Giovanni Casini, Thomas Meyer, and Ivan Varzinczak

Centre for Artificial Intelligence Research
CSIR Meraka Institute and UKZN, South Africa
{abritz,gcasini,tmeyer,ivarzinczak}@csir.co.za

Abstract. We extend the Description Logic \mathcal{ALC} with preferential role restrictions as class constructs, and argue that preferential universal restriction represents a defeasible version of standard universal restriction. The resulting DL is more expressive without adding to the complexity of TBox reasoning. We present a tableau system to compute TBox entailment, show that this notion of entailment is not sufficient when adding ABoxes, and refine entailment to deal adequately with ABox reasoning.

1 Introduction

Description logics (DLs) have been extended with features to express defeasibility in a number of ways, one of which is to incorporate preferential reasoning into the semantics [5, 10, 14, 20]. Recently, Britz, et al. [4] obtained a representation result for preferential and rational extensions of DLs linking previous syntactic approaches [10, 14] to semantic proposals for preferential extensions to DLs [5, 18, 20]. This established a foundation for the study of preferential and ranked entailment in DLs with a clear and intuitively appealing semantics supported by sound and complete reasoning support.

Different preferential extensions to DLs do not all share the same aim, and hence also do not share a semantics. One such aim is the representation of defeasible subsumption statements, which semantically translate to set inclusions admitting classical counter-examples. The focus there is therefore on defeasible statements of the form $C \sqsubseteq D$, read “ C s are usually D s” or “normally C s are D s”, leaving open the option for C s that are, in a sense, exceptional not to be instances of D . There are, however, a number of other aspects of defeasibility at the object level besides that of defeasible subsumption [1, 2, 14]. The common aim of these approaches is the introduction of some aspect of defeasibility, rather than non-monotonicity, with the latter rather emerging as a desired property of the resulting entailment relation in consequence of the introduction of the former.

Here we make a case for defeasible universal restrictions, in which a class description of the form $\forall r.C$ may be too strong, calling for a weaker version thereof which is defeasible in the sense that it admits classical counter-examples. As an example, consider the class description $\text{Lawyer} \sqcap \forall \text{hasClient.PayingClient}$, intended to capture the class of all private practice lawyers who only handle the cases of paying clients. This class description may be too strong, calling for a weaker class description of lawyers who normally defend only paying clients, but

who may exceptionally take on *pro bono* work. This leads to the introduction of defeasible universal restrictions. For example, $\text{Lawyer} \sqcap \forall \text{hasClient.PayingClient}$ can be used to describe the class of all lawyers having only paying clients, yet allowing for relatively exceptional role fillers to the `hasClient` role.

Dually, a class description of the form $\exists r.C$ may be too lenient, calling for a strengthening of the existential restriction construct that discounts exceptional role fillers. For example, $\exists \text{hasClient.PayingClient}$ describes the class of individuals having at least one paying client, whereas a description of the class of individuals whose normal clientèle includes at least one paying client requires a stricter version of existential restriction, written $\exists \text{hasClient.PayingClient}$. This notion can incidentally also be generalized to number restrictions.

The remainder of the paper is structured as follows: We present some background on preferential DL semantics in Section 2. In Section 3 we introduce preferential versions of universal and existential role restrictions, and define their semantics. We present a tableau system for \mathcal{ALC} TBox entailment with the added role restrictions in Section 4, and conclude in Section 6 with a few remarks on related and future work.

2 Preferential semantics for description logics

We assume the reader to be familiar with Description Logics, and follow standard notation. In this section we outline the preferential semantics for DLs obtained by enriching standard DL interpretations with an ordering on the elements in the domain. The intuition on which this approach is based is simple and natural, and extends similar work done for the propositional case [17, 19], and also more recently for description and modal logics [5–7, 9].

Informally, the semantics is based on the idea that objects of the domain can be ordered according to their degree of *normality* [3] or *typicality* [15, 2]. We do not require that objects intrinsically possess certain features that render some objects more normal than others. Rather, the intention is to provide a framework in which to express all conceivable ways in which objects, with their associated properties and relationships with other objects, can be ordered, in the same way that the class of all DL standard interpretations constitute a framework representing all conceivable ways of representing the properties of objects and their relationships with other objects. The knowledge base at hand therefore imposes constraints on the allowed orderings on objects in preferential DL interpretations in the same way as it imposes constraints on the allowed extensions of classes and roles in standard DL interpretations.

Definition 1 (Preferential Interpretation). *A preferential interpretation is a structure $\mathcal{P} = \langle \Delta^{\mathcal{P}}, \cdot^{\mathcal{P}}, \prec_{\mathcal{P}} \rangle$, where $\langle \Delta^{\mathcal{P}}, \cdot^{\mathcal{P}} \rangle$ is a DL interpretation (which we denote by $\mathcal{I}_{\mathcal{P}}$ and refer to as the standard interpretation associated with \mathcal{P}), and $\prec_{\mathcal{P}}$ is a strict partial order on $\Delta^{\mathcal{P}}$ (i.e., $\prec_{\mathcal{P}}$ is irreflexive and transitive) that is well-founded.¹*

¹ Observe that well-foundedness is a stricter condition to impose than the *smoothness* condition used for modelling defeasible subsumption [4].

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

Lemma 1. *Let $\mathcal{P} = \langle \Delta^{\mathcal{P}}, \cdot^{\mathcal{P}}, \prec_{\mathcal{P}} \rangle$ be a preferential interpretation. For every α , $\mathcal{P} \Vdash \alpha$ if and only if $\mathcal{I}_{\mathcal{P}} \Vdash \alpha$.*

Example 1. Let $\mathbf{N}_{\mathcal{C}} = \{A, B\}$ and let $\mathbf{N}_{\mathcal{R}} = \{r\}$. Figure 1 below depicts the preferential interpretation $\mathcal{P} = \langle \Delta^{\mathcal{P}}, \cdot^{\mathcal{P}}, \prec_{\mathcal{P}} \rangle$, where $\Delta^{\mathcal{P}} = \{x_i \mid 1 \leq i \leq 5\}$, $A^{\mathcal{P}} = \{x_1, x_2, x_3\}$, $B^{\mathcal{P}} = \{x_2, x_3, x_4\}$, $r^{\mathcal{P}} = \{(x_1, x_2), (x_2, x_3), (x_3, x_2), (x_1, x_4), (x_4, x_5), (x_5, x_4)\}$, which is represented by the solid arrows in the picture, and $\prec_{\mathcal{P}}$ is the transitive closure of $\{(x_1, x_2), (x_1, x_3), (x_2, x_4), (x_3, x_4), (x_4, x_5)\}$, i.e., of the relation represented by the dashed arrows in the picture. (Note the direction of the dashed arrows, pointing from less to more preferred objects, with more preferred objects lower in the order.)




Fig. 1. A preferential interpretation.

3 Enriching DLs with preferential role restrictions

The preferential DL interpretations presented above have been used elsewhere to define a defeasible subsumption relation, and also to define preferential and ranked entailment relations on defeasible DL knowledge bases [4]. Our present purpose is however to use it in the definition of a defeasible universal restriction class construct with a preferential semantics, analogous to the defeasible modalities defined by Britz, et al. [9]. We also define its dual class construct for existential restrictions. We show that preferential universal restriction introduces an aspect of defeasibility to the base class construct of universal restriction, while its dual introduces an aspect of strictness to the base class construct of existential restriction.

Definition 2. *Let $\mathcal{P} = \langle \Delta^{\mathcal{P}}, \cdot^{\mathcal{P}}, \prec_{\mathcal{P}} \rangle$ be a preferential interpretation. Given a role name r and a class description C , the truth conditions for defeasible universal restriction $\forall r.C$ and strict existential restriction $\exists r.C$ are given by:*

- $(\forall r.C)^{\mathcal{P}} := \{x \in \Delta^{\mathcal{P}} \mid \min_{\prec_{\mathcal{P}}} r^{\mathcal{P}}(x) \subseteq C^{\mathcal{P}}\}$.

- $(\exists r.C)^{\mathcal{P}} := \{x \in \Delta^{\mathcal{P}} \mid \min_{\prec_{\mathcal{P}}} r^{\mathcal{P}}(x) \cap C^{\mathcal{P}} \neq \emptyset\}$.

With $\tilde{\mathcal{L}}$ we denote the extension of \mathcal{L} obtained by adding \forall and \exists to the concept constructors of \mathcal{ALC} .

We say that $C \in \tilde{\mathcal{L}}$ is *preferentially satisfiable* if and only if there is a preferential interpretation \mathcal{P} such that $C^{\mathcal{P}} \neq \emptyset$. Preferential satisfiability of C with respect to a TBox \mathcal{T} is defined in the usual way. If $C, D \in \tilde{\mathcal{L}}$, then satisfaction of subsumption statements of the form $C \sqsubseteq D$ is just as before.

Definition 3. *Given a TBox \mathcal{T} and a subsumption statement α (built up from $\tilde{\mathcal{L}}$), we say that \mathcal{T} preferentially entails α , denoted $\mathcal{T} \models \alpha$, if and only if for every preferential interpretation \mathcal{P} , $\mathcal{P} \models \mathcal{K}$ implies $\mathcal{P} \models \alpha$.*

Lemma 2. *Let \mathcal{T} be a TBox and let $C, D \in \tilde{\mathcal{L}}$. $\mathcal{T} \models C \sqsubseteq D$ if and only if $C \sqcap \neg D$ is preferentially unsatisfiable with respect to \mathcal{T} .*

Definition 3 yields a monotonic entailment relation (in the sense that, if $\mathcal{T} \models \alpha$ then we also have that $\mathcal{T} \cup \{\beta\} \models \alpha$ for any subsumption statement β) with associated Tarskian consequence relation [9]. This raises the question of the nature and role of non-monotonicity in preferential role restrictions.

Non-monotonicity is often conflated in this way with defeasibility in the AI literature, probably because there *is* a connection between non-monotonic entailment relations and the intended application of such logics in defeasible reasoning. While not much harm may be done in that context, a more careful analysis is required when studying defeasibility of functions, operators or connectives in a language. We therefore briefly digress to recall some basic definitions surrounding these notions.

Definition 4. *Given $n + 1$ partially ordered sets of objects $\langle S_i, \leq_i \rangle$, $0 \leq i \leq n$, an n -ary function $f : \prod_{i=0}^{n-1} S_i \rightarrow S_n$ is monotone increasing in all its arguments on S_n if the following holds:*

$$\text{If } x_i \leq y_i \text{ for } 0 \leq i < n \text{ then } f(x_0, \dots, x_{n-1}) \leq f(y_0, \dots, y_{n-1}).$$

It is then easy to see that (the semantic interpretations of) the class constructs \sqcap , \sqcup and \exists induce monotone increasing binary functions on $\langle \mathcal{P}(\Delta), \sqsubseteq \rangle$, the class construct \neg induces a monotone decreasing unary function, and the class constructs \forall , \forall and \exists induce binary functions that are neither monotone increasing nor monotone decreasing, i.e., they are non-monotonic. We may alternatively take a more proof-theoretic approach and observe that \sqcap , \sqcup and \exists induce monotone increasing functions on the class subsumption hierarchy $\langle \mathcal{L}, \sqsubseteq \rangle$, \neg induces a monotone decreasing function and \forall , \forall and \exists induce non-monotonic functions on $\langle \mathcal{L}, \sqsubseteq \rangle$. Although we cannot express $r \sqsubseteq s$ in \mathcal{ALC} , semantically \forall remains a non-monotonic class construct, and in more expressive DLs with role hierarchies this can also be expressed syntactically, in that it does not in general follow from $r \sqsubseteq s$ and $C \sqsubseteq D$ that $\forall r.C \sqsubseteq \forall s.D$.

Defeasible reasoning dates back to Aristotle's analysis of dialectics, and relates to argument forms that seem compelling but are not classically valid.

As mentioned above, a more restrictive view of defeasible reasoning is often taken, conflating it with non-monotonic reasoning. However, demonstrating non-monotonicity to prove defeasibility is not always accurate. For example, as we showed above standard universal restriction is non-monotonic, but there is no reason why it should be regarded as a defeasible class construct.

Informally, a defeasible relation on a set is one which admits classical counter-examples. The defeasibility of the relation does not refer to its non-monotonicity, but rather to the possibility of defeat by counter-example. A requirement of a defeasible function is therefore that it is more *tolerant* than some base function in the following sense:

Definition 5. *Given $n + 1$ partially ordered sets of objects $\langle S_i, \leq_i \rangle$, $0 \leq i \leq n$, an n -ary function $f : \prod_{i=0}^{n-1} S_i \rightarrow S_n$ is tolerant with respect to an n -ary function $f' : \prod_{i=0}^{n-1} S_i \rightarrow S_n$ if the following holds:*

$$\begin{aligned} f'(x_1, \dots, x_n) &\leq f(x_1, \dots, x_n), \text{ for all } x_i \in S_i, 0 \leq i < n. \\ f'(x_1, \dots, x_n) &\neq f(x_1, \dots, x_n), \text{ for some } x_i \in S_i, 0 \leq i < n. \end{aligned}$$

The obvious examples relevant to the content of this paper are \forall , which is tolerant with respect to \forall , and \exists , which is tolerant with respect to \exists . Our claim is not that tolerance as defined above corresponds precisely to defeasibility (and could therefore be used as a definition of defeasibility). A simple illustration of this point is that classical disjunction is tolerant with respect to conjunction, but it does not seem to make sense to consider disjunction as being defeasible with respect to conjunction.

We contend that \forall may be interpreted as defeasible universal restriction: Informally, `PrivateLawyer` \equiv `Lawyer` \sqcap `\forall hasClient.PayingClient` defines the class of private lawyers as the set of objects that are lawyers, and whose typical clients are the ones who pay. The defeasibility resides in the class construct of universal quantification \forall , rather than in any of the class or role descriptions involved, or in the subsumption relation. Thus we are not stating that, normally, private lawyers only have paying clients, but rather that the normal clientèle of private lawyers are restricted to paying clients. This is made more precise in the example below.

Definition 6. *A TBox \mathcal{T} is said to be preferentially coherent if, for every $A \in \mathbf{N}_{\mathcal{E}}$, there is a preferential model \mathcal{P} of \mathcal{T} s.t. $A^{\mathcal{P}} \neq \emptyset$.²*

Example 2. Let \mathcal{T} be a TBox containing the following two statements:

`PrivateLawyer` \equiv `Lawyer` \sqcap `\forall hasClient.PayingClient`,
`CommunityLawyer` \sqsubseteq `PrivateLawyer` \sqcap `\exists hasClient. \neg PayingClient`.

It is easy to show that \mathcal{T} is preferentially coherent. Informally, this means it is possible for community lawyers, by virtue of being private lawyers, to normally have paying clients, but to have some non-paying clients at the same

² Preferential coherence is a generalisation of classical coherence as defined by Schlöbach, et al. [21].

time. Furthermore, it can be verified that \mathcal{T} preferentially entails the statement $\text{CommunityLawyer} \sqsubseteq \exists \text{hasClient.PayingClient}$. Informally, because we know that all community lawyers have some clients (albeit non-paying clients), it must be the case that they all have paying clients as well. This follows from the fact that community lawyers are also private lawyers.

4 Tableau system for preferential role restrictions

In this section we present a simple tableau-based proof procedure for reasoning with preferential role restrictions. Our tableau calculus is based on standard modal tableaux with labeled formulae and explicit accessibility relations [16]. (Our exposition here follows that given by Britz and Varzinczak [8] in the modal case, which is based on those by Castilho et al. [11, 12].)

As usually done in the DL community, subsumptions of the form $C \sqsubseteq D$, which are equivalent to $\top \sqsubseteq \neg C \sqcup D$, are treated internally by the tableau as concepts of the form $\neg C \sqcup D$. In that respect, it is enough to define a tableau system that checks only for concept satisfiability.

Definition 7. *If $n \in \mathbb{N}$ and $C \in \tilde{\mathcal{L}}$, then $n :: C$ is a labeled concept.*

In a labeled concept $n :: C$, the natural number n is the *label*. (As we shall see, informally, the idea is that the label stands for some object in a DL interpretation.)

Definition 8. *A skeleton is a function $\Sigma : \mathbb{N}_{\mathcal{R}} \rightarrow \mathcal{P}(\mathbb{N} \times \mathbb{N})$.*

That is, a skeleton maps role names in the language to binary relations on labels.

Our tableau system also makes use of an auxiliary structure of which the intention is to build a preference relation on objects of the domain:

Definition 9. *A preference relation \prec is a binary relation on \mathbb{N} .*

As we shall see below, like Σ , \prec is built cumulatively through successive applications of the tableau rules we shall introduce.

Definition 10. *A branch is a tuple $\langle \mathcal{S}, \Sigma, \prec \rangle$, where \mathcal{S} is a set of labeled concepts, Σ is a skeleton and \prec is a preference relation.*

Definition 11. *A tableau rule is a rule of the form:*

$$\rho \frac{\mathcal{N} ; \Gamma}{\mathcal{D}_1 ; \Gamma'_1 \mid \dots \mid \mathcal{D}_k ; \Gamma'_k}$$

where $\mathcal{N}; \Gamma$ is the numerator and $\mathcal{D}_1 ; \Gamma'_1 \mid \dots \mid \mathcal{D}_k ; \Gamma'_k$ is the denominator.

Given a rule ρ , \mathcal{N} represents one or more labeled concepts, called the *main concept* of the rule, separated by ‘;’. Γ stands for any additional *condition* (on Σ or \prec) that must be satisfied for the rule to be applicable (see below). In the

denominator, each \mathcal{D}_i , $1 \leq i \leq k$, has one or more labeled concepts, whereas each Γ'_i is a condition to be satisfied *after* the application of the rule (e.g. structural changes in the skeleton Σ or in the relation \prec). The symbol ‘|’ indicates the occurrence of a *split* in the branch, characteristic of the so-called *don't-know* non-deterministic rules.

Figure 2 below presents the set of tableau rules for $\tilde{\mathcal{L}}$. In the rules we abbreviate $(n, n') \in \Sigma(i)$ as $n \xrightarrow{i} n'$, and $n' \in \Sigma(i)(n)$ as $n' \in \Sigma_i(n)$. Finally, with n'^* , n''^* , \dots we denote labels that have not been used before. We say that a rule ρ is *applicable* to a branch $\langle \mathcal{S}, \Sigma, \prec \rangle$ if and only if \mathcal{S} contains an instance of the main concept of ρ and the conditions Γ of ρ are satisfied by Σ and \prec .

$$\begin{array}{l}
 (\perp) \frac{n :: C, n :: \neg C}{n :: \perp} \quad (\neg) \frac{n :: \neg \neg C}{n :: C} \quad (\cap) \frac{n :: C \cap D}{n :: C, n :: D} \quad (\sqcup) \frac{n :: \neg(C \cap D)}{n :: \neg C \mid n :: \neg D} \\
 (\forall) \frac{n :: \forall r_i.C ; n \xrightarrow{i} n', n' \in \min_{\prec} \Sigma_i(n)}{n' :: C} \quad (\exists) \frac{n :: \neg \forall r_i.C}{n'^* :: \neg C ; n \xrightarrow{i} n'^*, n'^* \in \min_{\prec} \Sigma_i(n)} \\
 (\forall) \frac{n :: \forall r_i.C ; n \xrightarrow{i} n'}{n' :: C} \quad (\exists) \frac{n :: \neg \forall r_i.C}{n'^* :: \neg C ; \Gamma'_1 \mid n'^* :: \neg C ; \Gamma'_2}, \text{ where} \\
 \Gamma'_1 = \{n \xrightarrow{i} n'^*, n'^* \in \min_{\prec} \Sigma_i(n)\} \text{ and} \\
 \Gamma'_2 = \{n \xrightarrow{i} n'^*, n \xrightarrow{i} n''^*, n''^* \prec n'^*, n''^* \in \min_{\prec} \Sigma_i(n)\}
 \end{array}$$

Fig. 2. Tableau rules for $\tilde{\mathcal{L}}$.

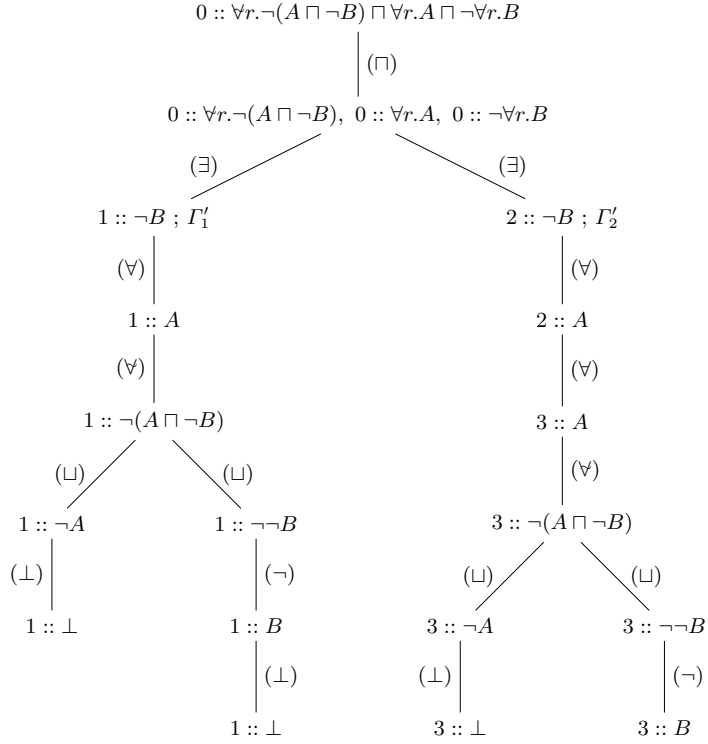
The Boolean rules together with (\forall) are as usual and need no explanation. Rule (\forall) propagates concepts in the scope of a defeasible universal restriction to the most preferred (with respect to \prec) of all successor nodes. Rule (\exists) creates a *preferred* (minimal) successor node with the corresponding labeled concept as content. Rule (\exists) replaces the standard rule for \exists -concepts with a don't-know non-deterministic version thereof and requires a more thorough explanation. When creating a new successor node, there are two possibilities: Either *(i)* it is minimal (with respect to \prec) amongst all successor nodes, in which case the result is the same as that of applying Rule (\exists) , or *(ii)* it is not minimal, in which case there must be a most preferred successor node that is more preferred (with respect to \prec) than the newly created one. (This splitting is of the same nature as that in the (\sqcup) -rule, i.e., it fits the purpose of a proof by cases.)

Definition 12. A tableau \mathcal{T} for $C \in \tilde{\mathcal{L}}$ is the limit of a sequence $\mathcal{T}^0, \dots, \mathcal{T}^n, \dots$ of sets of branches where the initial $\mathcal{T}^0 = \{\{\{0 :: C\}, \emptyset, \emptyset\}\}$ and every \mathcal{T}^{i+1} is obtained from \mathcal{T}^i by the application of one of the rules in Figure 2 to some branch $\langle \mathcal{S}, \Sigma, \prec \rangle \in \mathcal{T}^i$. Such a limit is denoted \mathcal{T}^∞ .

We make the so-called *fairness assumption*: Any rule that *can* be applied *will* eventually be applied, i.e., the order of rule applications is not relevant. We say a tableau is *saturated* if no rule is applicable to any of its branches.

Definition 13. A branch $\langle \mathcal{S}, \Sigma, \prec \rangle$ is closed if and only if $n :: \perp \in \mathcal{S}$ for some n . A saturated tableau \mathcal{T} for $C \in \tilde{\mathcal{L}}$ is closed if and only if all its branches are closed. (If \mathcal{T} is not closed, then we say that it is an open tableau.)

For an example, consider the concept $C = \exists r.(A \sqcap \neg B) \sqcup \exists r.\neg A \sqcup \forall r.B$. Figure 3 depicts the (open) tableau for $\neg C = \forall r.\neg(A \sqcap \neg B) \sqcap \forall r.A \sqcap \neg \forall r.B$.



$\Gamma'_1 = \text{add}(0, 1) \text{ to } \Sigma \text{ and } 1 \text{ to } \min_{\prec} \Sigma(0)$; $\Gamma'_2 = \text{add}(0, 2) \text{ and } (0, 3) \text{ to } \Sigma, (3, 2) \text{ to } \prec \text{ and } 3 \text{ to } \min_{\prec} \Sigma(0)$

Fig. 3. Visualization of an open tableau for a satisfiable concept.

From the open tableau in Figure 3 one can extract the preferential interpretation \mathcal{P} as depicted in Figure 4. (In Figure 4 the understanding is that $3 \prec 2$ and that 0 is *incomparable* with respect to \prec to the other objects in the domain.)

We are now ready to state the main result of this section. (The proof of Theorem 1 is analogous to that by Britz and Varzinczak in the modal case [8] and we do not state it here.)




Fig. 4. Preferential interpretation \mathcal{P} constructed from the tableau in Figure 3.

Theorem 1. $C \in \tilde{\mathcal{L}}$ is preferentially satisfiable if and only if there is an open (saturated) tableau for C .

It can easily be checked that in the construction of the tableau there is only a finite number of distinct states since every concept generated by the application of a rule is a subconcept of the original one. Therefore we have a decision procedure for our enriched description logic.

It is well-known that satisfiability checking for the description logic \mathcal{ALC} is PSPACE-complete. It is not that hard to see that the addition of \forall and \exists to the concept language does not affect the space complexity of the resulting tableaux. To see why, if the concept at the root of the tableau is C , and $|C| = m$, i.e., m is the number of symbols occurring in C , then the space requirement for each label is at most $O(m)$. Since there exists a saturated tableau with depth at most $O(m^2)$, the total space requirement is $O(m^3)$. In summary, in spite of the additional expressivity brought in by the introduction of preferential role restrictions, we remain in the same complexity class as that of the logic we started off with.

Finally, the tableau system we have just introduced checks only for concept satisfiability and therefore no TBox information is assumed. From the perspective of knowledge representation and reasoning it becomes important to check for concept satisfiability with respect to background knowledge provided in the form of a TBox and (possibly) an ABox. Fortunately our tableau calculus can easily be adapted to take care of this need. For instance, satisfiability with respect to a TBox can be achieved with the addition of the following rule [13]:

$$(TB) \frac{n :: E}{n :: E \sqcap \prod_{C \sqsubseteq D \in \mathcal{T}} (\neg C \sqcup D)}$$

The rule (TB) is the only one that does not have the subconcept property, but it is not hard to see that it does not affect decidability of the method. Complexitywise, since TBox entailment in \mathcal{ALC} is EXPTIME [13], and given our discussion above, one can expect entailment with preferential role restrictions to remain in EXPTIME, in particular if we require rule (TB) to be applied exactly once per node, which avoids an extra exponential blow up. A thorough analysis of complexity issues, as well as a study of proof strategies and optimizations, are beyond the scope of the present paper and we leave these for future work.

5 Reasoning with ABoxes

In the previous section we showed how to perform preferential TBox reasoning, and provided a tableau system for computing preferential TBox entailment. In this section we investigate what happens when we include ABoxes. Let $N_{\mathcal{I}}$ be a finite set of *individual names*. An \mathcal{ALC} ABox \mathcal{A} is a finite set of assertions of the form $C(a)$ where $a \in N_{\mathcal{I}}$ and $C \in \tilde{\mathcal{L}}$. We extend the preferential interpretation function $\cdot^{\mathcal{P}}$ of Definition 1 to include mappings from individual names to elements of the domain $\Delta^{\mathcal{P}}$ in the standard way: for every $a \in N_{\mathcal{I}}$, $a^{\mathcal{P}} \in \Delta^{\mathcal{P}}$. The satisfaction of (extended) ABox assertions is then defined as usual:

- For a preferential interpretation \mathcal{P} , and for $C \in \tilde{\mathcal{L}}$ and $a \in N_{\mathcal{I}}$, $\mathcal{P} \models C(a)$ if and only if $a^{\mathcal{P}} \in C^{\mathcal{P}}$.

An *extended \mathcal{ALC} knowledge base* \mathcal{K} is a tuple $\langle \mathcal{T}, \mathcal{A} \rangle$ where \mathcal{T} is an extended TBox and \mathcal{A} is an extended ABox. Henceforth we omit the word “extended” when referring to extended \mathcal{ALC} knowledge bases, TBoxes and ABoxes. We sometimes abuse notation by referring to $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ as the set $\mathcal{T} \cup \mathcal{A}$.

Given a knowledge base \mathcal{K} , the *preferential models* of \mathcal{K} is the set of preferential interpretations $\mathcal{PM}(\mathcal{K}) := \{\mathcal{P} \mid \mathcal{P} \models \alpha \text{ for every } \alpha \in \mathcal{K}\}$. \mathcal{K} is *preferentially satisfiable* if it has a preferential model. Definition 3 then applies without change also to preferential entailment from knowledge bases.

Example 3. Let \mathcal{T} be the TBox given in Example 2 and consider the ABox $\mathcal{A} = \{\text{PrivateLawyer}(\text{sam}), \exists \text{hasClient}.\top(\text{sam})\}$. The knowledge base $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ preferentially entails the ABox statement $\exists \text{hasClient}.\text{PayingClient}(\text{sam})$. That is, if we know that Sam the private lawyer has a client, we can conclude that he has a paying client. In fact, in line with the reasoning exhibited in Example 2, if we replace the statement $\exists \text{hasClient}.\top(\text{sam})$ in \mathcal{A} with the stronger statement $\exists \text{hasClient}.\neg \text{PayingClient}(\text{sam})$, the resulting knowledge base is preferentially satisfiable, and preferentially entails $\exists \text{hasClient}.\text{PayingClient}(\text{sam})$. In other words, given that Sam the private lawyer has a non-paying client, we can also conclude that he has a paying client.

At first glance, preferential entailment for knowledge bases may seem to provide appropriate results. However, closer inspection reveals that preferential entailment for knowledge bases sometimes gives counterintuitive results. Consider the knowledge base $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ where $\mathcal{T} = \{\text{Lawyer} \sqsubseteq \forall \text{hasClient}.\text{PayingClient}\}$ and $\mathcal{A} = \{\text{Lawyer}(\text{sam}), \text{hasClient}(\text{sam}, \text{peter})\}$. From this we would like to (defeasibly) conclude the statement $\text{PayingClient}(\text{peter})$. But is easy to see that preferential entailment does not sanction this conclusion. The issue is that we need individuals to be as typical as is allowed by the knowledge. For example, the reason why $\text{PayingClient}(\text{peter})$ is not entailed by \mathcal{K} is because there are some preferential models of \mathcal{K} in which Peter is not a paying client, and is therefore not one of Sam’s normal clients.

To rectify this, we introduce an ordering on the preferential models of a knowledge base. Intuitively, a preferential model \mathcal{P}_1 of \mathcal{K} is at least as preferred

(at least as low down in the ordering) as \mathcal{P}_2 if (i) it agrees with \mathcal{P}_2 everywhere except on the denotation of individual names, and (ii) all names in \mathcal{P}_1 denote objects that are at least as low down in the ordering as in \mathcal{P}_2 .

Definition 14. Given a knowledge base \mathcal{K} , we define the binary relation $\ll_{\mathcal{K}}$ on $\mathcal{PM}(\mathcal{K})$ as follows: $\mathcal{P}_1 \ll_{\mathcal{K}} \mathcal{P}_2$ if and only if $\Delta^{\mathcal{P}_1} = \Delta^{\mathcal{P}_2}$, $C^{\mathcal{P}_1} = C^{\mathcal{P}_2}$ for every $C \in \tilde{\mathcal{L}}$, $r^{\mathcal{P}_1} = r^{\mathcal{P}_2}$ for every $r \in \mathbf{N}_{\mathcal{D}}$, and $a^{\mathcal{P}_1} \preceq_{\mathcal{P}_1} a^{\mathcal{P}_2}$ for every $a \in \mathbf{N}_{\mathcal{I}}$.³

It is easily verified that $\ll_{\mathcal{K}}$ is a weak partial order (i.e., it is reflexive, anti-symmetric and transitive). Moreover, since all preferential interpretations are well-founded, it follows that $\ll_{\mathcal{K}}$ is also well-founded. Let $\min_{\ll_{\mathcal{K}}} \mathcal{PM}(\mathcal{K}) = \{\mathcal{P} \in \mathcal{PM}(\mathcal{K}) \mid \mathcal{P}' \not\ll_{\mathcal{K}} \mathcal{P} \text{ for every } \mathcal{P}' \in \mathcal{PM}(\mathcal{K}) \text{ s.t. } \mathcal{P}' \neq \mathcal{P}\}$.

Definition 15. Let \mathcal{K} be a knowledge base and α a (TBox or ABox) statement. α is said to be minimally preferentially entailed by \mathcal{K} , written $\mathcal{K} \models_{\ll} \alpha$, if and only if $\mathcal{P} \Vdash \alpha$ for every $\mathcal{P} \in \min_{\ll} \mathcal{PM}(\mathcal{K})$.

It can now be verified that the statement `PayingClient(peter)` is minimally preferentially entailed by the knowledge base in Example 3. To see that this is a defeasible conclusion, note that if we add the statement $\neg \text{PayingClient}(\text{peter})$ to the knowledge base in Example 3, the statement `PayingClient(peter)` is not minimally preferentially entailed by the new knowledge base (and the new knowledge base is preferentially satisfiable).

6 Concluding Remarks

We conclude with a comment on the expression of defeasible subsumption statements [4, 5, 7] in terms of defeasible universal restrictions. Given $C, D \in \mathcal{L}$, a statement of the form $C \sqsubseteq D$ is a *defeasible subsumption statement* and is read “usually C is subsumed by D ”. The connective \sqsubseteq is meant to be a defeasible counterpart of \sqsubseteq . A preferential interpretation \mathcal{P} satisfies a defeasible subsumption statement $C \sqsubseteq D$, written $\mathcal{P} \Vdash C \sqsubseteq D$, if and only if $\min_{\prec_{\mathcal{P}}} (C^{\mathcal{P}}) \subseteq D^{\mathcal{P}}$.

For every concept $C \in \mathcal{L}$ and preferential interpretation \mathcal{P} , we can define in \mathcal{P} its left cylindrification r_C in the following way:

$$r_C := \{\langle x, y \rangle \mid x \in \Delta^{\mathcal{P}} \text{ and } y \in C^{\mathcal{P}}\}.$$

Practically, r_C is the largest role r definable in a preferential model \mathcal{P} such that $\mathcal{P} \Vdash \top \sqsubseteq \forall r.C$. It follows that, for every $C, D \in \mathcal{L}$, $\mathcal{P} \Vdash C \sqsubseteq D$ if and only if $\mathcal{P} \Vdash \top \sqsubseteq \forall r_C.D$.

The implications of this connection between defeasible subsumption and defeasible universal restriction requires further investigation. Other future work include extension of the tableau procedure to minimal preferential entailment.

³ Note that $\preceq_{\mathcal{P}_1}$ is the weak partial order obtained from $\prec_{\mathcal{P}_1}$.

References

1. Bonatti, P., Lutz, C., Wolter, F.: The complexity of circumscription in description logic. *Journal of Artificial Intelligence Research* 35, 717–773 (2009)
2. Booth, R., Meyer, T., Varzinczak, I.: PTL: A propositional typicality logic. In: *Proc. JELIA*. pp. 107–119. No. 7519 in LNCS, Springer (2012)
3. Boutilier, C.: Conditional logics of normality: A modal approach. *Artificial Intelligence* 68(1), 87–154 (1994)
4. Britz, K., Casini, G., Meyer, T., Varzinczak, I.: Ordered Interpretations and Entailment for Defeasible Description Logics. Tech. rep., CAIR, CSIR Meraka and UKZN, South Africa (2013), <http://tinyurl.com/cydd6yy>
5. Britz, K., Heidema, J., Meyer, T.: Semantic preferential subsumption. In: *Proc. KR*. pp. 476–484. AAAI Press/MIT Press (2008)
6. Britz, K., Meyer, T., Varzinczak, I.: Preferential reasoning for modal logics. *Electronic Notes in Theoretical Computer Science* 278, 55–69 (2011)
7. Britz, K., Meyer, T., Varzinczak, I.: Semantic foundation for preferential description logics. In: *Proc. Australasian Joint Conference on AI*. pp. 491–500. No. 7106 in LNAI, Springer (2011)
8. Britz, K., Varzinczak, I.: Defeasible modes of inference: A preferential perspective. In: *Proc. NMR* (2012)
9. Britz, K., Varzinczak, I.: Defeasible modalities. In: *Proc. TARK*. pp. 49–60 (2013)
10. Casini, G., Straccia, U.: Rational closure for defeasible description logics. In: *Proc. JELIA*. pp. 77–90. No. 6341 in LNCS, Springer-Verlag (2010)
11. Castilho, M., Gasquet, O., Herzig, A.: Formalizing action and change in modal logic I: the frame problem. *Journal of Logic and Computation* 9(5), 701–735 (1999)
12. Castilho, M., Herzig, A., Varzinczak, I.: It depends on the context! A decidable logic of actions and plans based on a ternary dependence relation. In: *Proc. NMR* (2002)
13. Donini, F., Massacci, F.: EXPTIME tableaux for \mathcal{ALC} . *Artificial Intelligence* 124, 87–138 (2000)
14. Giordano, L., Olivetti, N., Gliozzi, V., Pozzato, G.: $\mathcal{ALC} + T$: a preferential extension of description logics. *Fundamenta Informaticae* 96(3), 341–372 (2009)
15. Giordano, L., Olivetti, N., Gliozzi, V., Pozzato, G.: A minimal model semantics for nonmonotonic reasoning. In: *Proc. JELIA*. pp. 228–241. No. 7519 in LNCS, Springer (2012)
16. Goré, R.: Tableau methods for modal and temporal logics. In: D’Agostino, M., Gabbay, D., Hähnle, R., Posegga, J. (eds.) *Handbook of Tableau Methods*, pp. 297–396. Kluwer Academic Publishers (1999)
17. Kraus, S., Lehmann, D., Magidor, M.: Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence* 44, 167–207 (1990)
18. Lehmann, D., Magidor, M.: Preferential logics: the predicate calculus case. In: *Proc. TARK*. pp. 57–72 (1990)
19. Lehmann, D., Magidor, M.: What does a conditional knowledge base entail? *Artificial Intelligence* 55, 1–60 (1992)
20. Quantz, J.: A preference semantics for defaults in terminological logics. In: *Proc. KR*. pp. 294–305 (1992)
21. Schlöblich, S., Cornet, R.: Non-standard reasoning services for the debugging of description logic terminologies. In: *Proc. IJCAI*. pp. 355–362 (2003)