Adding Default Attributes to $\mathcal{EL}^{++}$

Piero A. Bonatti and **Marco Faella** and **Luigi Sauro**

Dip. di Scienze Fisiche, Università di Napoli Federico II

Abstract

The research on low-complexity nonmonotonic description logics recently identified a fragment of $\mathcal{EL}^{\perp}$, supporting defeasible inheritance with overriding, where reasoning can be carried out in polynomial time. We contribute to that framework by supporting more axiom schemata and all the concepts of $\mathcal{EL}^{++}$ without increasing asymptotic complexity.

Introduction

Description logics (DLs) are an essential component of the semantic web. Most prominent ontologies are formulated with DLs; the OWL standard is founded on DLs; DLs have also been used as security policy languages, capable of expressing access constraints based on semantic metadata, see for example (Uszok et al. 2003; Finin et al. 2008; Zhang et al. 2009; Kolovski, Hendler, and Parsia 2007). In order to address the efficiency requirements posed by the size of the semantic web, researchers have identified description logics whose reasoning tasks can be carried out in polynomial time, such as *DL-lite* (Calvanese et al. 2005) (that provides a foundation to RDFS) and the $\mathcal{EL}$ family of logics (Baader 2003; Baader, Brandt, and Lutz 2005), that are general enough to express important biomedical ontologies of practical interest such as SNOMED and (most of) GALEN. These fragments constitute the foundation of two of the OWL2 dialects. All of these description logics are monotonic, therefore they cannot address some application requirements that call for (nonmonotonic) forms of inheritance supporting overriding and exceptions. A first such requirement stems from the biomedical domain, where it is common to define prototypical entities whose default properties can be later refined in subclasses. Since DLs currently do not support any such mechanism, some efforts have been devoted to workarounds that—however—do not provide a general answer to this issue (Rector 2004; Stevens et al. 2007). Similarly, policy languages usually support authorization inheritance with exceptions for selected user subgroups and/or object subclasses; moreover, nonmonotonic inferences are needed to model conflict resolution strategies and common default policies such as *open* and *closed* policies (cf. (Bonatti and Samarati 2003)).

A recent work (Bonatti, Faella, and Sauro 2010) provides a first answer to the above needs by identifying a circumscribed version of $\mathcal{EL}^{\perp}$ where concepts can be associated to default attributes that can be overridden in subconcepts; reasoning is in PTIME. In this paper, we extend that framework by supporting more axiom schemata and all the concept constructors of $\mathcal{EL}^{++}$. We show that, for this purpose, the notion of *conflict safety* adopted in (Bonatti, Faella, and Sauro 2010) to achieve tractability needs to be adapted. Moreover, we add limited support to variable concept names and prove that unrestricted variable concepts and role composition make reasoning intractable.

Paper organization: After a preliminary section, where $\mathcal{EL}^{++}$ and nonmonotonic $\mathcal{EL}^{\perp}$ are recalled, we show the limitations of the original notion of conflict safety that appear when the representation language is generalized. Then, in the fourth section, we provide a solution, prove tractability results, and discuss some further extensions that affect tractability. Two sections devoted, respectively, to related works and a final discussion conclude the paper.

Preliminaries

We assume the reader to be familiar with the syntax and semantics of monotonic Description Logics. We refer to (Baader et al. 2003, Chap. 2) for details and notation. The sets of *concept names*, *role names*, and *individual names* are denoted by N_C, N_R, and N_I, respectively. $\mathcal{EL}^{++}$ supports also a set N_F of *concrete features*. By *predicate* we mean any member of $N_C \cup N_R \cup N_F$. Hereafter, letters A, B, H, and K range over N_C, P and R range over N_R, a, b, c range over N_I, and f over N_F. Letters C, D range over concepts.

Syntax and semantics of the logic $\mathcal{EL}^{++}$ (Baader, Brandt, and Lutz 2005) are shown in Figure 1.[1] Recall that interpretations are pairs $\mathcal{I} = \langle \Delta^{\mathcal{I}}, \cdot^{\mathcal{I}} \rangle$ where $\Delta^{\mathcal{I}}$ is the interpretation domain and function $\cdot^{\mathcal{I}}$ maps concept names and role names on subsets of $\Delta^{\mathcal{I}}$ and $\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$, respectively. Concrete features are interpreted by means of an additional fixed structure $\mathcal{D} = \langle \Delta^{\mathcal{D}}, \mathcal{P}^{\mathcal{D}} \rangle$ called *concrete domain*, where $\Delta^{\mathcal{D}}$ is a domain and $\mathcal{P}^{\mathcal{D}}$ a set of predicates. Each n-ary $p \in \mathcal{P}^{\mathcal{D}}$ is interpreted as a relation $p^{\mathcal{D}} \subseteq (\Delta^{\mathcal{D}})^n$. For all $f \in N_F$ and all interpretations $\mathcal{I}$, $f^{\mathcal{I}}$ is a partial function from $\Delta^{\mathcal{I}}$ to $\Delta^{\mathcal{D}}$. In

---

[1] So-called *assertions* and *ABoxes* are not included because they can be simulated with nominals.

Name	Syntax	Semantics
top	$\top$	$\top^{\mathcal{I}} = \Delta^{\mathcal{I}}$
bottom	$\bot$	$\bot^{\mathcal{I}} = \emptyset$
conjunction	$C \sqcap D$	$C^{\mathcal{I}} \cap D^{\mathcal{I}}$
existential restriction	$\exists R.C$	$\{d \in \Delta^{\mathcal{I}} \mid \exists (d,e) \in R^{\mathcal{I}} : e \in C^{\mathcal{I}}\}$
nominal	$\{a\}$	$\{a^{\mathcal{I}}\}$
concrete domains	$p(f_1, \ldots, f_n)$	$\{x \in \Delta^{\mathcal{I}} \mid f_i(x) \text{ is defined } (1 \leq i \leq n)$ and $\langle f_1(x), \ldots, f_n(x)\rangle \in p^{\mathcal{D}}\}$
GCI	$C \sqsubseteq D$	$C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$
RI	$r_1 \circ \cdots \circ r_n \sqsubseteq r$	$r_1^{\mathcal{I}} \circ \cdots \circ r_n^{\mathcal{I}} \subseteq r^{\mathcal{I}}$

Figure 1: Syntax and semantics of $\mathcal{EL}^{++}$.

general, multiple concrete domains can be supported. They must be *P-admissible* to make reasoning tractable; here we omit the definition to meet space limitations—see (Baader, Brandt, and Lutz 2005) for all details.

A *(strong) knowledge base* is a finite set of *general concept inclusions* (GCIs) and *role inclusions* (RIs); $\mathcal{I}$ is a *model* of a strong knowledge base $\mathcal{KB}$ iff $\mathcal{I}$ satisfies all the elements of $\mathcal{KB}$ as specified in Fig. 1. We write $C \sqsubseteq_{\mathcal{KB}} D$ iff for all models $\mathcal{I}$ of $\mathcal{KB}$, $\mathcal{I}$ satisfies $C \sqsubseteq D$.

Next, we recall the nomonotonic framework introduced in (Bonatti, Faella, and Sauro 2010). A *defeasible inclusion* (DI) is an expression $A \sqsubseteq_n C$ whose intended meaning is: *A's elements are normally in C*. For all $\delta = A \sqsubseteq_n C$, let $\mathsf{pre}(\delta) = A$ and $\mathsf{con}(\delta) = C$.

Example 1 The sentences: *"in humans, the heart is usually located on the left-hand side of the body; in humans with situs inversus, the heart is located on the right-hand side of the body"* (Rector 2004; Stevens et al. 2007) can be formalized with the $\mathcal{EL}^{++}$ axioms and DI:

> Human $\sqsubseteq_n$ ∃has_heart.∃has_position.Left ;
> Situs_Inversus $\sqsubseteq$ ∃has_heart.∃has_position.Right ;
> ∃has_heart.∃has_position.Left $\sqcap$
>   ∃has_heart.∃has_position.Right $\sqsubseteq \bot$. ∎

A *defeasible knowledge base* (DKB) is the disjoint union $\mathcal{KB} = \mathcal{KB}_S \cup \mathcal{KB}_D$ of a strong knowledge base $\mathcal{KB}_S$ and a set of DIs $\mathcal{KB}_D$. For the sake of readability, $C \sqsubseteq_{\mathcal{KB}_S} D$ will be abbreviated to $C \sqsubseteq_{\mathcal{KB}} D$.

DIs are prioritized using the following *specificity* relation $\prec$: For all DIs δ_1 and δ_2, let

$\delta_1 \prec \delta_2$ iff $\mathsf{pre}(\delta_1) \sqsubseteq_{\mathcal{KB}} \mathsf{pre}(\delta_2)$ and $\mathsf{pre}(\delta_2) \not\sqsubseteq_{\mathcal{KB}} \mathsf{pre}(\delta_1)$.

With specificity, in case of conflicts the DIs associated to subconcepts override those associated to their superconcepts, by analogy with the inheritance mechanism of object-oriented languages.

Example 2 The access control policy: *"Normally users cannot read project files; staff can read project files; blacklisted staff is not granted any access"* can be encoded with:

> Staff $\sqsubseteq$ User
> Blacklisted $\sqsubseteq$ Staff
> UserRequest $\equiv$ ∃subj.User $\sqcap$ ∃target.Proj $\sqcap$ ∃op.Read
> StaffRequest $\equiv$ ∃subj.Staff $\sqcap$ ∃target.Proj $\sqcap$ ∃op.Read
> UserRequest $\sqsubseteq_n$ ∃decision.Deny
> StaffRequest $\sqsubseteq_n$ ∃decision.Grant
> ∃subject.Blacklisted $\sqsubseteq$ ∃decision.Deny
> ∃decision.Grant $\sqcap$ ∃decision.Deny $\sqsubseteq \bot$.

(The expressions of the form $C \equiv D$ abbreviate the two inclusions $C \sqsubseteq D$ and $D \sqsubseteq C$.) Staff members cannot simultaneously satisfy the two defeasible inclusions (due to the last inclusion above). With specificity, the second defeasible inclusion *overrides* the first one and yields the intuitive inference that non-blacklisted staff members are indeed allowed to access project files. ∎

The semantics of DKBs specializes Circumscription. It consists in maximizing the set of individuals satisfying the DIs in $\mathcal{KB}_D$, resolving conflicts with specificity whenever possible. During maximization, the extension of some predicates may vary, while others are fixed. The choice of (Bonatti, Faella, and Sauro 2010) consists in fixing concept names and letting roles vary. The rationale is the following: The main goal is adding default attributes to existing classical ontologies. An invasive mechanism that can extensively change the semantics of concepts is undesirable in this context.[2] By fixing concept names, atomic concepts preserve their possible extensions (i.e., their classical semantics); only attributes (i.e. role types and values, and concrete feature values) may change. In general, however, the set of concept names $\mathsf{N_C}$ can be arbitrarily partitioned into two sets F and V containing fixed and varying predicates, respectively. Fixed predicates retain their classical semantics, while varying predicates can be affected by nonmonotonic inferences.

Now for the formal definitions. The set F, the DIs $\mathcal{KB}_D$, and their ordering $\prec$ induce a strict preference ordering over interpretations, defined below. Roughly speaking, $\mathcal{I}$ is preferred to $\mathcal{J}$ if some DIs are satisfied by more individuals in $\mathcal{I}$ than in $\mathcal{J}$, possibly at the cost of satisfying less lower-priority DIs. Formally, for all $\delta = (C \sqsubseteq_n D)$ and all interpretations $\mathcal{I}$ let the set of individuals *satisfying* δ be:

$$\mathsf{sat}_{\mathcal{I}}(\delta) = \{x \in \Delta^{\mathcal{I}} \mid x \notin C^{\mathcal{I}} \text{ or } x \in D^{\mathcal{I}}\} .$$

Definition 3 Let $\mathcal{KB}$ be a DKB. For all interpretations $\mathcal{I}$ and $\mathcal{J}$, and all $F \subseteq \mathsf{N_C}$, let $\mathcal{I} <_{\mathcal{KB},F} \mathcal{J}$ iff:

1. $\Delta^{\mathcal{I}} = \Delta^{\mathcal{J}}$;
2. $a^{\mathcal{I}} = a^{\mathcal{J}}$, for all $a \in \mathsf{N_I}$;
3. $A^{\mathcal{I}} = A^{\mathcal{J}}$, for all $A \in F$;
4. for all $\delta \in \mathcal{KB}_D$, if $\mathsf{sat}_{\mathcal{I}}(\delta) \not\supseteq \mathsf{sat}_{\mathcal{J}}(\delta)$ then there exists $\delta' \in \mathcal{KB}_D$ such that $\delta' \prec \delta$ and $\mathsf{sat}_{\mathcal{I}}(\delta') \supset \mathsf{sat}_{\mathcal{J}}(\delta')$;
5. there exists a $\delta \in \mathcal{KB}_D$ such that $\mathsf{sat}_{\mathcal{I}}(\delta) \supset \mathsf{sat}_{\mathcal{J}}(\delta)$.

The subscript $\mathcal{KB}$ will be omitted when clear from context. Now a *model* of a DKB can be defined as a maximally preferred model of its strong (i.e. classical) part.

Definition 4 (Model) Let $F \subseteq \mathsf{N_C}$. An interpretation $\mathcal{I}$ is a *model* of $\mathsf{Circ}_F(\mathcal{KB})$ iff $\mathcal{I}$ is a (classical) model of $\mathcal{KB}_S$ and for all models $\mathcal{J}$ of $\mathcal{KB}_S$, $\mathcal{J} \not<_F \mathcal{I}$.

In order to enhance readability, $<_{\mathsf{var}}$ and $\mathsf{Circ}_{\mathsf{var}}$ stand for $<_\emptyset$ and $\mathsf{Circ}_\emptyset$, respectively; $<_{\mathsf{fix}}$ and $\mathsf{Circ}_{\mathsf{fix}}$ stand respectively for $<_{\mathsf{N_C}}$ and $\mathsf{Circ}_{\mathsf{N_C}}$.

In (Bonatti, Faella, and Sauro 2010) we introduced a special reasoning task for retrieving the default properties of

---

[2]For example, a varying concept name may become unsatisfiable when its default properties override some default property of a super-concept.

any given concept A. This is not completely trivial, because in general a member of A may belong to some $B \sqsubseteq A$, whose specific default properties may override some of the prototypical properties of A. In order to avoid this problem (that results in incomplete answers), the members of A's subclasses are implicitly removed by rewriting subsumption $A \sqsubseteq C$ to $CWA_{\mathcal{KB}}(A) \sqsubseteq C$, where $CWA_{\mathcal{KB}}(A) = A \sqcap \prod\{\neg B \mid B \in \mathsf{N_C}$ and $A \not\sqsubseteq_{\mathcal{KB}} B\} \sqcap \prod\{\neg\{a\} \mid a \in \mathsf{N_I}$ and $A \not\sqsubseteq_{\mathcal{KB}} \{a\}\}$. In (Bonatti, Faella, and Sauro 2010) we argued that this transformation preserves the intended default attributes of A. By

$$\mathsf{Circ}_F(\mathcal{KB}) \models_{\mathsf{cw}} A \sqsubseteq C$$

we mean that $CWA_{\mathcal{KB}}(A) \sqsubseteq C$ is satisfied by all the models of $\mathsf{Circ}_F(\mathcal{KB})$. We proved that $\models_{\mathsf{cw}}$ can be decided in polynomial time under the following assumptions:

1. $F = \mathsf{N_C}$ (i.e., $\mathsf{Circ_{fix}}$ is adopted);

2. the language is $\mathcal{EL}^\perp$, that is, the restriction of $\mathcal{EL}^{++}$ obtained by eliminating nominals, concrete domains, and role inclusions;

3. $\mathcal{KB}$ and subsumptions are restricted to instances of the inclusion schemata (1) below, where A and B range over atomic concepts;

$$
\begin{array}{llll}
A \sqsubseteq B & A_1 \sqsubseteq \exists R.A_2 & A_1 \sqcap A_2 \sqsubseteq B \\
\exists P.A \sqsubseteq B & \exists P.A_1 \sqsubseteq \exists R.A_2 & \exists P.A_1 \sqcap \exists R.A_2 \sqsubseteq \perp
\end{array}
\quad (1)
$$

4. $\mathcal{KB}$ is *conflict safe*; roughly speaking, this means that whenever two DIs $A_1 \sqsubseteq_n C_1$ and $A_2 \sqsubseteq_n C_2$ have incomparable priorities and $C_1 \sqcap C_2$ is unsatisfiable w.r.t. $\mathcal{KB}$, either A_1 and A_2 are disjoint, or there must be a more specific DI $A_1 \sqcap A_2 \sqsubseteq_n C_3$ that blocks at least one of the two conflicting DIs. In Bonatti, Faella, and Sauro we proved also that if $\mathcal{KB}$ is not required to be conflict safe then reasoning becomes coNP-hard.

The restriction to conflict safe KBs in practice means that ontology authors should define consistent prototypical entities, which is reasonable in typical application domains. If unresolved conflicts arise from multiple inheritance, then knowledge engineers are responsible for removing conflicts with specific DIs. Fortunately, conflict safety checking as well as a simple automatic repair strategy are in PTIME (see (Bonatti, Faella, and Sauro 2010) for more details).

Limitations of Conflict Safety

Conflict safety suffices to make reasoning tractable only if $\mathcal{KB}$ conforms to the restricted axiom schemata in (1). If more general axioms are allowed, then conflicts may arise that are not detected by the definition of conflict safety. As a consequence, deciding whether $\mathsf{Circ_{fix}}(\mathcal{KB}) \models_{\mathsf{cw}} A \sqsubseteq D$ holds is coNP-hard in the generalized framework. This can be proved by reducing SAT to $\models_{\mathsf{cw}}$. For each clause c_i in the SAT instance, introduce two roles C_i and $\bar{C}_i$. Intuitively, the meaning of $\exists C_i$ and $\exists \bar{C}_i$ is: c_i is/is not satisfied, respectively. For each propositional symbol p_j introduce roles P_j, $\bar{P}_j$. Intuitively, $\exists P_j$ and $\exists \bar{P}_j$ represent the truth of literals p_j and $\neg p_j$, respectively. An additional role P^* will be illustrated later. Next, we need three concept names B_0, B_1 and B_2

such that $B_0 \sqsubseteq B_1$ and $B_1 \sqsubseteq B_2$, and a role $\bar{F}$. Intuitively, $\exists \bar{F}$ represents the falsity of the set of clauses. Axiomatize clauses by adding:

$$\exists P_j \sqsubseteq \exists C_i, \quad \exists \bar{P}_k \sqsubseteq \exists C_i,$$

for all disjuncts p_j and $\neg p_k$ in c_i. The space of possible truth assignments is generated by the following inclusions:

$$
\begin{array}{llll}
B_1 \sqsubseteq_n \exists P_j & (2) & B_0 \sqsubseteq \exists P^* & (4) \\
B_1 \sqsubseteq_n \exists \bar{P}_j & (3) & \exists P_j \sqcap \exists \bar{P}_j \sqcap \exists P^* \sqsubseteq \perp. & (5)
\end{array}
$$

DIs (2) and (3) have the same priority and "block" each other in B_0, due to (4) and (5); so (2) and (3) induce a complete truth assignment. Then we introduce a defeasible inclusion with lower priority: $B_2 \sqsubseteq_n \exists \bar{C}_i$. This defeasible inclusion "assumes" that c_i is not satisfied. The other axioms may defeat this assumption (if the selected truth assignment entails $\exists C_i$) due to the disjointness axiom: $\exists C_i \sqcap \exists \bar{C}_i \sqsubseteq \perp$.

Finally, add the inclusions $\exists \bar{C}_i \sqsubseteq \exists \bar{F}$ to say that the set of clauses is not satisfied when at least one of the clauses is false. Now let $\mathcal{KB}$ denote the above set of inclusions. It can be proved that the given set of clauses is unsatisfiable iff:

$$\mathsf{Circ_{fix}}(\mathcal{KB}) \models_{\mathsf{cw}} B_0 \sqsubseteq \exists \bar{F}.$$

It follows that reasoning is coNP-hard although this $\mathcal{KB}$ is conflict safe (because the right-hand sides of the incomparable DIs (2) and (3) are mutually consistent w.r.t. $\mathcal{KB}_S$). The above reduction relies on an inclusion that does not conform to (1), namely, (5), that could also be indirectly encoded if (1) were extended with the more general schema $\exists P_1 \sqcap \exists P_2 \sqsubseteq \exists R$, as the reader may easily verify.

A similar problem arises with concrete domains. Take for example the P-admissible domain $\langle \mathbb{Q}, \{=_1, +_1\} \rangle$, where $=_1(f)$ is satisfied iff feature f has value 1, and $+_1(f, g)$ is satisfied iff $f + 1 = g$ (Baader, Brandt, and Lutz 2005). The above reduction can be easily adapted by uniformly replacing $\exists P_j$ with $=_1(f_j)$ and $\exists \bar{P}_j$ with $=_1(\bar{f}_j)$; moreover, (4) should be replaced by $B_0 \sqsubseteq +_1(f_j, \bar{f}_j)$, so that $=_1(f_j)$ and $=_1(\bar{f}_j)$ cannot be simultaneously satisfied in B_0 (the same result was obtained by (5) in the previous reduction).

Summarizing, we have:

Theorem 5 *Let $\mathcal{KB}$ range over conflict safe DKBs. Deciding $\mathsf{Circ_{fix}}(\mathcal{KB}) \models_{\mathsf{cw}} A \sqsubseteq \exists P$ is coNP-hard if $\mathcal{KB}$ may contain any of the following features:*

1. *instances of the schema $\exists P_1.A_1 \sqcap \exists P_2.A_2 \sqcap \exists P_3.A_3 \sqsubseteq \perp$;*

2. *instances of the schema $\exists P_1.A_1 \sqcap \exists P_2.A_2 \sqsubseteq \exists R.A$;*

3. *concrete domains.*

Consequently, we are going to strengthen conflict safety in order to restore tractability in the generalized framework.

Strong Conflict Safety

In this section we generalize the old approach to a "standard" fragment of $\mathsf{Circ_{fix}}(\mathcal{EL}^{++})$. For the sake of simplicity, we assume that $\mathcal{KB}$'s strong concept inclusions are instances of the following schemata that generalize (1), where F ranges over the (fixed) concepts $\mathsf{N_C} \cup \{\{a\} \mid a \in \mathsf{N_I}\} \cup \{\perp, \top\}$, metavariables V, V_1, V_2, V_3 range over the

(variable) concepts of the form $p(f_1, \ldots, f_n)$ and $\exists R.B$, and B ranges over $\mathsf{N_C} \cup \{\top\}$:

$$F \sqsubseteq V \qquad V \sqsubseteq F \qquad V_1 \sqsubseteq V_2 \qquad V_1 \sqcap V_2 \sqsubseteq V_3 \, .$$

$\mathcal{KB}$ may contain $\mathcal{EL}^{++}$ role inclusions with no role composition. These schemata constitute a normal form for classical $\mathcal{EL}^{++}$ without role composition and a normal form for the fragment of $\mathsf{Circ_{fix}}(\mathcal{EL}^{++})$ with no role composition and no quantifier nesting (that increase complexity, see below).

A DI $A \sqsubseteq_n \exists R.B$ is blocked when B is empty, because circumscription cannot affect fixed concepts. In real ontologies (especially biomedical ontologies), concept names are typically meant to be nonempty, as witnessed by the use of concept satisfiability checking as a debugging tool. In practice, nonemptiness assumptions need not be made explicit in a classical (monotonic) setting, while in circumscribed $\mathcal{EL}^{\perp}$ and $\mathcal{EL}^{++}$ their omission may improperly block the derivation of default attributes, cf. (Bonatti, Faella, and Sauro 2010). Due to space limitations, here we simply assume that all concept names B are explicitly declared to be nonempty (say, with inclusions like $\top \sqsubseteq \exists aux.B$, where aux is a fresh role). A more articulated approach will be discussed in an extended version of this paper.

We say that $\mathcal{KB}$ is *standard* if it satisfies the above assumptions (i.e., no quantifier nesting, no role composition, and nonempty concept names).

In order to characterize the DIs applicable to a concept name A, we collect the DIs that apply to A's super-concepts H and are not blocked by more specific DIs applicable to A:

$$\mathsf{nonblocked}(A,H) = \big\{ \delta \in \mathcal{KB}_D \mid \mathsf{pre}(\delta) \equiv_{\mathcal{KB}} H,$$
$$CWA_{\mathcal{KB}}(A) \sqcap \mathsf{con}(\delta) \sqcap \mathsf{inh}(A,H) \not\sqsubseteq_{\mathcal{KB}} \perp \big\},$$

where $\mathsf{inh}(A, H)$ collects the default properties that can be inherited from the concepts between A and H, excluding H:

$$\mathsf{inh}(A, H) = \bigsqcap \{ \mathsf{con}(\delta') \mid \delta' \in \mathsf{nonblocked}(A, B), B \in [A, H) \},$$

$$[A, H) = \{ B \in \mathsf{N_C} \mid A \sqsubseteq_{\mathcal{KB}} B, B \sqsubseteq_{\mathcal{KB}} H, H \not\sqsubseteq_{\mathcal{KB}} B \} \, .$$

Finally, the set of all properties that can be inherited by A is:

$$\mathsf{inh}(A) = \bigsqcap \{ \mathsf{con}(\delta) \mid \delta \in \mathsf{nonblocked}(A, H) \text{ and } A \sqsubseteq_{\mathcal{KB}} H \} \, .$$

Now, by analogy with the old notion of conflict safety, we require $A \sqcap \mathsf{inh}(A)$ to be satisfiable; in other words, the set of DIs that are not individually blocked by any group of more specific DIs should be simultaneously applicable:

Definition 6 A standard $\mathcal{KB}$ is *strongly conflict safe* iff for all A occurring in $\mathcal{KB}$ ($A \in \mathsf{N_C}$), $A \sqcap \mathsf{inh}(A)$ is satisfiable w.r.t. $\mathcal{KB}$.

In practice—as in the old framework—this means that conflicts that cannot be resolved by specificity must be explicitly removed with specific DIs. Strong conflict safety restores tractability if quantifiers are not nested. Let $depth(C)$ be the maximum quantifier nesting level of C:

Theorem 7 *If $A \sqcap \mathsf{inh}(A)$ is satisfiable with respect to a standard $\mathcal{KB}$ and $depth(C) = 0$, then $\mathsf{Circ_{fix}}(\mathcal{KB}) \models_{\mathsf{cw}} A \sqsubseteq C$ iff $A \sqcap \mathsf{inh}(A) \sqsubseteq_{\mathcal{KB}} C$.*

Proof. In the following, we say that an individual *actively satisfies* a DI δ, if it satisfies both $\mathsf{pre}(\delta)$ and $\mathsf{con}(\delta)$. Moreover, we set $\mathsf{nonblocked}(A) = \bigcup_{A \sqsubseteq_{\mathcal{KB}} H} \mathsf{nonblocked}(A, H)$.

[if] Let $\mathcal{I} \in \mathsf{Circ_{fix}}(\mathcal{KB})$ such that *(i)* $x \in A^{\mathcal{I}}$ and *(ii)* by definition of $\models_{\mathsf{cw}}$ and $CWA_{\mathcal{KB}}(A)$, if $x \in B^{\mathcal{I}}$ then $A \sqsubseteq_{\mathcal{KB}} B$. We prove that x actively satisfies a defeasible inclusion $\delta \in C^{\mathcal{I}}$ if and only if $\delta \in \mathsf{nonblocked}(A)$. As a consequence, $x \in C^{\mathcal{I}}$.

By contradiction, either x actively satisfies $\delta \notin \mathsf{nonblocked}(A)$ or x violates $\delta \in \mathsf{nonblocked}(A)$. We define an interpretation $\mathcal{I}'$ s.t. $\mathcal{I}' <_{\mathsf{fix}} \mathcal{I}$. We only redefine the roles and features involving x.

1. Let $\delta \in \mathsf{nonblocked}(A)$. Observe that $\mathsf{con}(\delta)$ can be a conjunction of concept names, existential restrictions $\exists R.B$ and concrete predicates $p(f_1, \ldots, f_n)$. For each existential restriction $\exists R.B$ in $\mathsf{con}(\delta)$, let $y \in B^{\mathcal{I}}$, we set $(x, y) \in R^{\mathcal{I}'}$. For each predicate $p(f_1, \ldots, f_n)$ in $\mathsf{con}(\delta)$, let $h_1, \ldots, h_n \in \Delta^{\mathcal{D}}$ be such that $(h_1, \ldots, h_n) \in p^{\mathcal{D}}$. Such concrete values exist because $\delta \in \mathsf{nonblocked}(A)$ and hence $p(f_1, \ldots, f_n)$ is satisfiable w.r.t. $\mathcal{KB}$. We set $f_j^{\mathcal{I}'}(x) = h_j$ for all $j = 1, \ldots, n$. In this way, x satisfies in $\mathcal{I}'$ all DIs in $\mathsf{nonblocked}(A)$.

2. Recursively, if there is a concept C s.t. $x \in C^{\mathcal{I}'}$ and $C \sqsubseteq_{\mathcal{KB}} \exists S.K$, then let $y \in K^{\mathcal{I}}$, we set $(x, y) \in S^{\mathcal{I}'}$. Additionally, if $C \sqsubseteq_{\mathcal{KB}} p(f_1, \ldots, f_n)$, let $h_1, \ldots, h_n \in \Delta^{\mathcal{D}}$ be such that $(h_1, \ldots, h_n) \in p^{\mathcal{D}}$, we set $f_j^{\mathcal{I}'}(x) = h_j$ for all $j = 1, \ldots, n$. Finally, if $R \sqsubseteq_{\mathcal{KB}} S$ and $(x, y) \in R^{\mathcal{I}'}$, we set $(x, y) \in S^{\mathcal{I}'}$. This rule helps ensuring that $\mathcal{I}$ is a classical model of $\mathcal{KB}$.

First, $\mathcal{I}'$ is a classical model of $\mathcal{KB}$ due to step 2 and the assumption that $A \sqcap \mathsf{inh}(A)$ is satisfiable.
Next, we show that $\mathcal{I}'$ is preferred to $\mathcal{I}$ w.r.t. $<_{\mathsf{fix}}$. Clearly, any DI $\delta \in \mathsf{nonblocked}(A)$ that x violates in $\mathcal{I}$ is strictly improved by $\mathcal{I}'$. Now, let $\delta \notin \mathsf{nonblocked}(A)$ be such that x actively satisfies it in $\mathcal{I}$, and let $H = \mathsf{pre}(\delta)$. We show that there is another DI with higher priority than δ that is strictly improved by $\mathcal{I}'$. Since x only satisfies A and its superclasses, we have $A \sqsubseteq_{\mathcal{KB}} H$. Since $\delta \notin \mathsf{nonblocked}(A)$, $CWA_{\mathcal{KB}}(A) \sqcap \mathsf{con}(\delta) \sqcap \mathsf{inh}(A, H)$ is unsatisfiable in $\mathcal{KB}$. Since $CWA_{\mathcal{KB}}(A) \sqcap \mathsf{con}(\delta)$ is satisfiable, $\mathsf{inh}(A, H)$ is not empty and there is $B \in [A, H)$ and $\delta' \in \mathsf{nonblocked}(A, B)$ such that x violates δ' in $\mathcal{I}$. Notice that δ' has a higher priority than δ because $B \sqsubseteq_{\mathcal{KB}} H$ and $H \not\sqsubseteq_{\mathcal{KB}} B$.

[only if] We prove the contrapositive. Assume that $A \sqcap \mathsf{inh}(A) \not\sqsubseteq_{\mathcal{KB}} C$. We define an interpretation $\mathcal{I}$ in $\mathsf{Circ_{fix}}(\mathcal{KB})$ that contains an individual x such that: *(i)* $x \in A^{\mathcal{I}}$, *(ii)* if $x \in B^{\mathcal{I}}$ then $A \sqsubseteq_{\mathcal{KB}} B$, and *(iii)* $x \notin C^{\mathcal{I}}$. We set $\Delta^{\mathcal{I}} = \{d_B \mid B \in \mathsf{N_C}\} \cup \{d_a \mid a \in \mathsf{N_I}\}$, where each individual d_B (resp., d_a) satisfies only the concept name B (resp., nominal a) and its superclasses. We only specify the roles and features involving x, as the others are not relevant.

1. Let $\delta \in \mathsf{nonblocked}(A)$. For each existential restriction $\exists R.B$ in $\mathsf{con}(\delta)$, we set $(x, d_B) \in R^{\mathcal{I}}$. For each predicate $p(f_1, \ldots, f_n)$ in $\mathsf{con}(\delta)$, let $h_1, \ldots, h_n \in \Delta^{\mathcal{D}}$ be such that $(h_1, \ldots, h_n) \in p^{\mathcal{D}}$. We set $f_j^{\mathcal{I}}(x) = h_j$ for all $j = 1, \ldots, n$.

2. Recursively, if there is a concept C s.t. $x \in C^{\mathcal{I}}$ and $C \sqsubseteq_{\mathcal{KB}} \exists S.K$, then we set $(x, d_K) \in S^{\mathcal{I}}$. Additionally, if $C \sqsubseteq_{\mathcal{KB}} p(f_1, \ldots, f_n)$, let $h_1, \ldots, h_n \in \Delta^{\mathcal{D}}$ be such that $(h_1, \ldots, h_n) \in p^{\mathcal{D}}$, we set $f_j^{\mathcal{I}}(x) = h_j$ for all $j = 1, \ldots, n$. Finally, if $R \sqsubseteq_{\mathcal{KB}} S$ and $(x, y) \in R^{\mathcal{I}}$, we set $(x, y) \in S^{\mathcal{I}}$.

It remains to show that $\mathcal{I} \in \mathsf{Circ}_{\mathsf{fix}}(\mathcal{KB})$. Due to step 2 and the fact that $A \sqcap \mathsf{inh}(A)$ is satisfiable w.r.t. $\mathcal{KB}$, $\mathcal{I}$ is a classical model of $\mathcal{KB}$. Moreover, assume by contradiction that there exists an interpretation $\mathcal{I}'$ that is a model of $\mathcal{KB}$ and such that $\mathcal{I}' <_{\mathsf{fix}} \mathcal{I}$. Since concept names are fixed, we have (i) $x \in A^{\mathcal{I}'}$, and (ii) if $x \in B^{\mathcal{I}'}$ then $A \sqsubseteq_{\mathcal{KB}} B$. Then, there is a DI δ that is actively satisfied by x in $\mathcal{I}'$, violated by x in $\mathcal{I}$, and such that no other DI with higher priority than δ has the same property. Observe that $\delta \notin \mathsf{nonblocked}(A)$. Let $H = \mathsf{pre}(\delta)$, since $x \in H^{\mathcal{I}'}$ and by (ii) above, we have $A \sqsubseteq_{\mathcal{KB}} H$. All DIs that have higher priority than δ and are satisfied by x in $\mathcal{I}$ are satisfied in $\mathcal{I}'$ as well. Formally, for all $B \in [A, H)$ and for all $\delta' \in \mathsf{nonblocked}(A, B)$, x actively satisfies δ' in $\mathcal{I}'$. By definition of $\mathsf{inh}(A, H)$, we obtain that $CWA_{\mathcal{KB}}(A) \sqcap \mathsf{con}(\delta) \sqcap \mathsf{inh}(A, H)$ is satisfiable w.r.t. $\mathcal{KB}$.

On the other hand, the only reason why $\delta \notin \mathsf{nonblocked}(A)$ is that the conjunction $CWA_{\mathcal{KB}}(A) \sqcap \mathsf{con}(\delta) \sqcap \mathsf{inh}(A, H)$ is not satisfiable w.r.t. $\mathcal{KB}$, which is a contradiction. ∎

Since classical $\mathcal{EL}^{++}$ subsumption is in PTIME, we immediately get:

Corollary 8 *If $\mathcal{KB}$ is strongly conflict safe and $depth(C) = 0$, then deciding $\mathsf{Circ}_{\mathsf{fix}}(\mathcal{KB}) \models_{\mathsf{cw}} A \sqsubseteq C$ is in PTIME.*

Examples: The reduction illustrated in the previous section is not strongly conflict safe: $\mathsf{inh}(B_0) = \exists P_j \sqcap \exists \bar{P}_j$ because none of the two DIs is blocked by any group of more specific DIs; consequently, by (5), $B_0 \sqcap \mathsf{inh}(B_0)$ is not satisfiable in $\mathcal{KB}$. On the contrary, it is easy to verify that the knowledge bases in examples 1 and 2 are strongly conflict safe. An interesting artificial example is:

$$A \sqsubseteq A_1, \quad A \sqsubseteq A_2, \quad A_2 \sqsubseteq B,$$

$$A_1 \sqsubseteq_n \exists R_1, \quad A_2 \sqsubseteq_n \exists R_2, \quad B \sqsubseteq_n \exists P.$$

First assume the additional inclusion $\exists R_1 \sqcap \exists R_2 \sqcap \exists P \sqsubseteq \bot$. This $\mathcal{KB}$ is not strongly conflict safe: $\mathsf{inh}(A) = \exists R_1 \sqcap \exists R_2 \sqcap \exists P$ (because no DI is blocked by any more specific DIs) therefore $\mathsf{inh}(A)$ is unsatisfiable in $\mathcal{KB}$. Indeed, for some individuals the first DI blocks the third one, for others the opposite happens, so $\mathsf{Circ}_{\mathsf{fix}}(\mathcal{KB})$ entails the disjunctive inclusion $A \sqsubseteq \exists R_1 \sqcup \exists P$. Such disjunctions in general increase complexity. Similarly, if the additional inclusion is $\exists R_1 \sqcap \exists P \sqsubseteq \bot$, there is an unresolved conflict across two diverging branches of the taxonomy and $\mathcal{KB}$ is not strongly conflict safe. Finally, if the additional inclusion is $\exists R_2 \sqcap \exists P \sqsubseteq \bot$, then $\mathcal{KB}$ is strongly conflict safe: the conflict between the second and third DIs is resolved by specificity and $\mathsf{Circ}_{\mathsf{fix}}(\mathcal{KB})$ entails $A \sqsubseteq \exists R_1 \sqcap \exists R_2$. Any conflict arising across different branches must be explicitly resolved. For instance, consider again the conflict arising from $\exists R_1 \sqcap \exists P \sqsubseteq \bot$. It may be resolved in favor of the first DI with $A \sqsubseteq_n \exists R_1$ (which overrides the third DI). Symmetrically, it can be resolved in favor of the third DI with

$A \sqsubseteq_n \exists P$. As a (mechanizable) default strategy, it is possible to block all conflicting DIs in A by introducing a fresh role F and the inclusions

$$A \sqsubseteq_n \exists F, \quad \exists F \sqcap \exists R_1 \sqsubseteq \bot, \quad \exists F \sqcap \exists P \sqsubseteq \bot. \quad \blacksquare$$

Unfortunately, the hypothesis $depth(C) = 0$ is essential to the correctness of the above results. In the following we show a reduction of SAT that exploits nested existentials in the query. Each pair of literals p_j and $\neg p_j$ is encoded by two (fixed) concept names P_j and $\bar{P}_j$, respectively, constrained to be disjoint by $P_j \sqcap \bar{P}_j \sqsubseteq \bot$. For all clauses c_i, the concept name $\bar{C}_i$ represents c_i's falsity:

$$\bigsqcap \{\bar{P}_j \mid p_j \in c_i\} \cup \{P_k \mid \neg p_k \in c_i\} \sqsubseteq \bar{C}_i.$$

Two more concepts $F, \bar{F}$—representing the truth values of the given clause set—are axiomatized by:

$$\bar{C}_i \sqsubseteq \bar{F}, \quad F \sqcap \bar{F} \sqsubseteq \bot.$$

Now, if the clause set is unsatisfiable, then F can be satisfied by an individual d only if d belongs neither to P_j nor to $\bar{P}_j$, for some j. Such "undefined" truth values are detected with:

$$P_j \sqcap \exists U_j \sqsubseteq \bot, \quad \bar{P}_j \sqcap \exists U_j \sqsubseteq \bot, \quad \exists U_j \sqsubseteq \exists U, \quad F \sqsubseteq_n \exists U_j.$$

Let $\mathcal{KB}$ be the union of $A_0 \equiv \exists R.F$ and the above set of inclusions. It is not hard to see that $\mathsf{Circ}_{\mathsf{fix}}(\mathcal{KB}) \models_{\mathsf{cw}} A_0 \sqsubseteq \exists R.\exists U$ iff the given clauses are unsatisfiable. Consequently, deciding $\models_{\mathsf{cw}}$ is coNP-hard in the presence of quantifier nesting, even if nesting is confined in the query and the maximum nesting level is one.

A similar negative result holds for $\mathsf{Circ}_{\mathsf{var}}$ (and hence Circ_F). It suffices to observe that every nested concept $\exists P.\exists R.A$ can be simulated by a non-nested concept $\exists P.B$, where B is a fresh variable concept defined by $B \equiv \exists R.A$.[3] This reduction can be further adapted to prove that role composition affects tractability. Summarizing:

Theorem 9 *Let $\mathcal{KB}$ range over strongly conflict safe DKBs. Deciding $\mathsf{Circ}_F(\mathcal{KB}) \models_{\mathsf{cw}} A \sqsubseteq C$ is coNP-hard if some of the following conditions may hold:*

1. $depth(C) > 0$ *(i.e., quantifiers can be nested);*
2. $F \subset \mathsf{N_C}$ *(i.e., there are variable concepts);*
3. $\mathcal{KB}$ *is extended with RIs with role composition.*

The expressiveness of variable concept names can be partially recovered. It is not hard to see that variable concept names can be allowed wherever an unqualified restriction $\exists P$ would be. Essentially, this means that the only constraint is: in each concept $\exists R.A$, A must be a fixed atom.

Related Work

The idea of maximizing the set of individuals satisfying a distinguished set of (possibly prioritized) axioms can be found in Defeasible Logic (Nute 1994; Antoniou et al. 2001) and Courteous Logic Programs (Grosof 1997). Reproducing the same semantics in a description logic framework is not trivial, e.g. Herbrand domains are infinite due to existential

---

[3]If B is fixed, then this construction changes $\mathcal{KB}$'s semantics: the equivalence fixes the (otherwise variable) expression $\exists R.A$.

quantification, so rule instantiation yields infinite sets, and decidability may be affected. The alternative is instantiating rules only on constants that explicitly occur in $\mathcal{KB}$; this means that default properties do not apply to the implicit (unnamed) individuals that exist due to existential quantification. These issues have been tackled (but not completely solved) in autoepistemic frameworks (Donini, Nardi, and Rosati 1997; 2002) using standard names as interpretation domains. These approaches, as well as those based on default logic (Baader and Hollunder 1995) and circumscription (Bonatti, Lutz, and Wolter 2009), in general range from NExpTime^{NP} to undecidability. A few restricted fragments based on circumscription lie within the second level of the polynomial hierarchy (Cadoli, Donini, and Schaerf 1990; Bonatti, Faella, and Sauro 2009). Similarly, the recent approaches (Giordano et al. 2009; Casini and Straccia 2010) are intractable. The only fragment in PTIME currently known is (Bonatti, Faella, and Sauro 2010), that here has been extended as explained in the previous sections.

Discussion and Conclusions

We proved that the extension of $\mathcal{EL}^{\perp}$ with default attributes introduced in (Bonatti, Faella, and Sauro 2010) can be generalized to $\mathcal{EL}^{++}$ without role composition, and to more general inclusion schemata without affecting tractability. For this purpose, conflict safety must be adapted in order to handle conflicts arising from groups of three or more inclusions. Moreover, the argument of existential quantifiers must always be a fixed concept. We proved that unrestricted variable predicates and quantifier nesting make reasoning coNP-hard. Note that strong conflict safety is *not always* more restrictive than conflict safety: the latter requires a "repairing" DI δ with $\mathsf{pre}(\delta) \equiv_{\mathcal{KB}} \mathsf{pre}(\delta_1) \sqcap \mathsf{pre}(\delta_2)$ whenever $\mathcal{KB}$ contains two incomparable conflicting defaults δ_1 and δ_2, while strong conflict safety allows the conflict to be resolved by any number of DIs whose left-hand side is not necessarily equivalent to $\mathsf{pre}(\delta_1) \sqcap \mathsf{pre}(\delta_2)$. One of the interesting features of the old framework is the ability of checking conflict safety and providing default automatic repair (by blocking all the conflicting DI pairs) in polynomial time. The same is possible in the new framework: Note that the sets $\mathsf{inh}(A)$ underlying strong conflict safety can be computed by solving a polynomial number of $\mathcal{EL}^{++}$ subsumptions (each of which is in PTIME). When $A \sqcap \mathsf{inh}(A)$ is unsatisfiable, the responsible DIs can be blocked by adding more DIs to A. The next steps will concern the support of explicit priorities, efficient implementations, and experimentations aimed at verifying the framework's scalability.

References

Antoniou, G.; Billington, D.; Governatori, G.; and Maher, M. J. 2001. Representation results for defeasible logic. *ACM Trans. Comput. Log.* 2(2):255–287.

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

Baader, F.; McGuiness, D. L.; Nardi, D.; and Patel-Schneider, P. 2003. *The Description Logic Handbook: Theory, implementation and applications*. Cambridge University Press.

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

Baader, F. 2003. The instance problem and the most specific concept in the description logic EL w.r.t. terminological cycles with descriptive semantics. In *Proc. of the 26th Annual German Conf. on AI, KI 2003*, volume 2821 of *LNCS*, 64–78. Springer.

Bonatti, P. A., and Samarati, P. 2003. Logics for authorization and security. In *Logics for Emerging Applications of Databases*, 277–323. Springer.

Bonatti, P. A.; Faella, M.; and Sauro, L. 2009. Defeasible inclusions in low-complexity DLs: Preliminary notes. In *IJCAI*, 696–701.

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

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

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

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

Casini, G., and Straccia, U. 2010. Rational closure for defeasible description logics. In *JELIA*, to appear.

Donini, F. M.; Nardi, D.; and Rosati, R. 1997. Autoepistemic description logics. In *IJCAI (1)*, 136–141.

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

Finin, T. W.; Joshi, A.; Kagal, L.; Niu, J.; Sandhu, R. S.; Winsborough, W. H.; and Thuraisingham, B. M. 2008. ROWLBAC: representing role based access control in OWL. In *SACMAT*, 73–82. ACM.

Giordano, L.; Gliozzi, V.; Olivetti, N.; and Pozzato, G. L. 2009. Prototypical reasoning with low complexity description logics: Preliminary results. In *LPNMR*, volume 5753 of *LNCS*, 430–436. Springer.

Grosof, B. N. 1997. Prioritized conflict handling for logic programs. In *ILPS*, 197–211.

Kolovski, V.; Hendler, J. A.; and Parsia, B. 2007. Analyzing web access control policies. In *WWW*, 677–686. ACM.

Nute, D. 1994. Defeasible logic. In *Handbook of logic in artificial intelligence and logic programming*, volume 3, 353–395.

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

Stevens, R.; Aranguren, M. E.; Wolstencroft, K.; Sattler, U.; Drummond, N.; Horridge, M.; and Rector, A. L. 2007. Using OWL to model biological knowledge. *International Journal of Man-Machine Studies* 65(7):583–594.

Uszok, A. *et al.* 2003. KAoS policy and domain services: Towards a description-logic approach to policy representation, deconfliction, and enforcement. In *4th IEEE Int. Workshop on Policies for Distributed Systems and Networks (POLICY)*, 93–96. IEEE Computer Soc.

Zhang, R.; Artale, A.; Giunchiglia, F.; and Crispo, B. 2009. Using description logics in relation based access control. In *Description Logics*, volume 477 of *CEUR Workshop Proc.* CEUR-WS.org.