Computing Role-depth Bounded Generalizations in the Description Logic $\mathcal{ELOR}$

Andreas Ecke[1][*], Rafael Peñaloza[1,2][**], and Anni-Yasmin Turhan[1][***]

[1] Institute for Theoretical Computer Science,
Technische Universität Dresden
[2] Center for Advancing Electronics Dresden
{ecke,penaloza,turhan}@tcs.inf.tu-dresden.de

Abstract. Description Logics (DLs) are a family of knowledge representation formalisms, that provides the theoretical basis for the standard web ontology language OWL. Generalization services like the least common subsumer (lcs) and the most specific concept (msc) are the basis of several ontology design methods, and form the core of similarity measures. For the DL $\mathcal{ELOR}$, which covers most of the OWL 2 EL profile, the lcs and msc need not exist in general, but they always exist if restricted to a given role-depth. We present algorithms that compute these role-depth bounded generalizations. Our method is easy to implement, as it is based on the polynomial-time completion algorithm for $\mathcal{ELOR}$.

1 Introduction

Description logics (DLs) are knowledge representation formalisms with formal and well-understood semantics [4]. They supply the foundation for the web ontology language OWL 2 standardized by the W3C [20]. Since then, DLs became more widely used for the representation of knowledge from several domains.

Each DL offers a set of concept constructors by which complex concepts can be built. These concepts describe categories from the application domain at hand. A DL knowledge base consists of two parts: the TBox captures the terminological knowledge about categories and relations, and the ABox captures the assertional knowledge, i.e., individual facts, from the application domain. Prominent inferences are *subsumption*, which determines subconcept relationships and *instance checking*, which tests for a given individual and concept whether the individual belongs to the concept.

The lightweight DL $\mathcal{EL}$ offers limited expressivity but allows for polynomial time reasoning [3]. These good computational properties are maintained by several extensions of $\mathcal{EL}$—most prominently by $\mathcal{EL}^{++}$, the DL underlying the OWL 2 EL profile [15], which allows for the use of nominals, i.e., singleton concepts, when building complex concept descriptions. The reasoning algorithms

for (fragments of) $\mathcal{EL}^{++}$ have been implemented in highly optimized reasoner systems such as jCEL [14] and ELK [10]. It is worth pointing out that the initial reasoning algorithm for handling nominals in $\mathcal{EL}$ [3] turned out to be incomplete, but a complete method has been recently devised in [11].

In this paper we describe methods for computing generalizations in the $\mathcal{EL}$-family with the help of the standard reasoning algorithms. We consider the following two inferences: The *least common subsumer* (lcs), which computes for a given set of concepts a new concept that subsumes the input concepts and is the least one w.r.t. subsumption; and the *most specific concept* which provides a concept which has a given individual as an instance and is the least one w.r.t. subsumption. Both inferences have been employed for several applications. Most prominently the lcs and the msc can be employed in the 'bottom-up approach' for generating TBoxes, where modellers can generate a new concept from picking ABox individuals that instantiate the desired concept and then generalizing this set into a single concept automatically–first by applying the msc to each of the individuals and then generalizing the obtained concepts by applying the lcs [5]. Other applications of the lcs and the msc include similarity measures [8, 6, 13], which are the core of ontology matching algorithms and more (see [7, 16]). In particular for large bio-medical ontologies the lcs can be used effectively to support construction and maintenance. Many of these bio-medical ontologies, notably SNOMED CT [19] and the FMA Ontology [18], are written in the $\mathcal{EL}$-family of lightweight DLs.

It is known that for concepts captured in a general TBox or even just a cyclic TBox, the lcs does not need to exist [1], since cycles cannot be captured in an $\mathcal{EL}$-concept. Therefore, an approximation has been introduced in [16], that limits the maximal nesting of quantifiers of the resulting concept descriptions. These so-called role-depth bounded lcs (k-lcs), can be computed for $\mathcal{EL}$ and for $\mathcal{EL}$ extended by role inclusions using completion sets produced by the subsumption algorithm [16, 9]. In this paper, we describe a subsumption algorithm for the DL $\mathcal{ELOR}$—building on the one for $\mathcal{ELO}$ ($\mathcal{EL}$ extended by nominals) from [11]. Our algorithm is given in terms of the completion algorithm in order to extend the methods for the k-lcs to $\mathcal{ELOR}$.

Recently, necessary and sufficient conditions for the existence of the lcs w.r.t. general $\mathcal{EL}$-TBoxes have been devised [21]. By the use of these conditions the bound k for which the role-depth bounded lcs and the lcs coincide can be determined, if the lcs exists; i.e., if such k is finite.

Similarly to the lcs, the msc does not need to exist, if the ABox [12] contain cycles. To obtain an approximative solution, the role-depth of the resulting concept can be limited as suggested in [12]. A computation algorithm for the role-depth bounded msc has been proposed in [17] for $\mathcal{EL}$. If nominals are allowed, the computation of the msc is trivial, since the msc of an individual a is simply the nominal that contains a (i.e., $\{a\}$). Thus, we consider the computation of the role-depth bounded msc in $\mathcal{EL}$ w.r.t. an $\mathcal{ELOR}$ knowledge base.

We introduce the basic notions of DL and the reasoning services considered in the next section. In Section 3 we give a completion-based classification algo-

	Syntax	Semantics
concept name	$A \;\; (A \in N_C)$	$A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$
top concept	$\top$	$\Delta^{\mathcal{I}}$
nominal	$\{a\} \;\; (a \in N_I)$	$\{a\}^{\mathcal{I}} = \{a^{\mathcal{I}}\}$
conjunction	$C \sqcap D$	$(C \sqcap D)^{\mathcal{I}} = C^{\mathcal{I}} \cap D^{\mathcal{I}}$
existential restriction	$\exists r.C \;\; (r \in N_R)$	$(\exists r.C)^{\mathcal{I}} = \{d \in \Delta^{\mathcal{I}} \mid \exists e.(d,e) \in r^{\mathcal{I}} \wedge e \in C^{\mathcal{I}}\}$
GCI	$C \sqsubseteq D$	$C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$
RIA	$r_1 \circ \cdots \circ r_n \sqsubseteq s$	$r_1^{\mathcal{I}} \circ \cdots \circ r_n^{\mathcal{I}} \subseteq s^{\mathcal{I}}$
Concept assertion	$C(a)$	$a^{\mathcal{I}} \in C^{\mathcal{I}}$
Role assertion	$r(a,b)$	$(a^{\mathcal{I}}, b^{\mathcal{I}}) \in r^{\mathcal{I}}$

Table 1. Concept constructors and TBox axioms for $\mathcal{ELOR}$.

rithm for $\mathcal{ELOR}$, which serves as a basis for the computation algorithms of the role-depth bounded lcs and msc presented subsequently. The paper ends with conclusions and future work.[3]

2 Preliminaries

$\mathcal{ELOR}$-*concepts* are built from mutually disjoint sets N_C of *concept names*, N_R of *role names* and N_I of *individual names* using the syntax rule:

$$C, D ::= \top \mid A \mid \{a\} \mid C \sqcap D \mid \exists r.C,$$

where $A \in N_C$, $r \in N_R$ and $a \in N_I$. The individuals appearing in concepts are also called *nominals*. The sub-logic of $\mathcal{ELOR}$ that does not allow for individuals in concepts is called $\mathcal{ELR}$.

As usual, the semantics of $\mathcal{ELOR}$-concepts is defined through interpretations. An *interpretation* $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ consists of an *interpretation domain* $\Delta^{\mathcal{I}}$ and an *interpretation function* $\cdot^{\mathcal{I}}$ that maps concept names A to subsets $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$ and role names to binary relations on the domain $\Delta^{\mathcal{I}}$. This function is extended to complex concepts as shown in the upper part of Table 1.

Concepts can be used to model notions from the application domain in the TBox. Given two concepts C and D, a *general concept inclusion axiom* (GCI) is of the form $C \sqsubseteq D$. We use $C \equiv D$ as an abbreviation for $C \sqsubseteq D$ and $D \sqsubseteq C$. Given the roles $r_1, \ldots, r_n$ and s, a *role inclusion axiom* (RIA) is an expression of the form $r_1 \circ \cdots \circ r_n \sqsubseteq s$. An $\mathcal{ELOR}$-TBox is a finite set of GCIs and RIAs. An interpretation is a *model* for a TBox $\mathcal{T}$ if it satisfies all GCIs and RIAs in $\mathcal{T}$, as shown in the middle part of Table 1. An $\mathcal{EL}$-TBox is an $\mathcal{ELR}$-TBox (i.e., without the nominal constructor) that does not contain any RIAs.

Knowledge about individual facts of the application domain can be captured by assertions. Let $a, b \in N_I$, $r \in N_R$ and C a concept, then $C(a)$ is a *concept*

---

[3] Because of space constraints, some proofs are deferred to the appendix of long version of this paper at `http://lat.inf.tu-dresden.de/research/papers.html`.

assertion and $r(a, b)$ a *role assertion*. An *ABox* $\mathcal{A}$ is a finite set of (concept or role) assertions. An interpretation is a *model* for an ABox $\mathcal{A}$ if it satisfies all concept and role assertions in $\mathcal{A}$, as shown in the lower part of Table 1.

A *knowledge base* (KB) $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ consists of a TBox $\mathcal{T}$ and an ABox $\mathcal{A}$. An interpretation is a model of $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ if it is a *model* of both $\mathcal{T}$ and $\mathcal{A}$. With $\mathsf{Sig}(\mathcal{T})$ we denote the *signature* of a TBox $\mathcal{T}$, i.e. the set of all concept names, role names, and individual names that appear in $\mathcal{T}$. By $\mathsf{Sig}(\mathcal{A})$ and $\mathsf{Sig}(\mathcal{K})$ we denote the analogous notions for ABoxes and KBs, respectively.

Important reasoning tasks considered for DLs are *subsumption* and *instance checking*. A concept C *is subsumed by* a concept D w.r.t. a TBox $\mathcal{T}$ (denoted $C \sqsubseteq_{\mathcal{T}} D$) if $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$ holds in all models $\mathcal{I}$ of $\mathcal{T}$. A concept C is *equivalent* to a concept D w.r.t. a TBox $\mathcal{T}$ (denoted $C \equiv_{\mathcal{T}} D$) if $C \sqsubseteq_{\mathcal{T}} D$ and $D \sqsubseteq_{\mathcal{T}} C$ hold. The reasoning service *classification* of a TBox $\mathcal{T}$ computes all subsumption relationships between the named concepts occurring in $\mathcal{T}$. A reasoning service dealing with a whole KB is *instance checking*. An individual a *is an instance of* a given concept C w.r.t. $\mathcal{K}$ (denoted $\mathcal{K} \models C(a)$) if $a^{\mathcal{I}} \in C^{\mathcal{I}}$ holds in all models $\mathcal{I}$ of $\mathcal{K}$. *ABox realization* computes, for every concept name in $\mathcal{K}$, the set of individuals from the ABox that belong to that concept. These reasoning problems can all be decided for $\mathcal{ELOR}$, and hence also in $\mathcal{EL}$, in polynomial time [3].

There are two central inferences discussed in this paper that compute generalizations. The first is called the *least common subsumer* (lcs); it computes, for two given concepts, a (possibly complex) concept that subsumes both input concepts and that is the least concept with this property w.r.t. subsumption. The second is called the *most specific concept* (msc), which computes for a given individual a the least concept w.r.t. subsumption that has a as an instance w.r.t. $\mathcal{K}$.

The lcs does not need to exist if computed w.r.t. general $\mathcal{EL}$-TBoxes, i.e., TBoxes that use complex concepts in the left-hand sides of GCIs, or even just cyclic TBoxes [2]. The reason is that the resulting concept cannot capture cycles. Thus, we follow here the idea from [16] and compute only approximations of the lcs and of the msc by limiting the nesting of quantifiers of the resulting concept.

The *role depth $rd(C)$* of a concept C denotes the maximal nesting depth of the existential quantifier in C. Sometimes it is convenient to write the resulting concept in a different DL than the one the inputs concepts are written in. Thus we distinguish a 'source DL' $\mathcal{L}_s$ and a 'target DL' $\mathcal{L}_t$. With these notions at hand, we can define the first generalization inference.

Definition 1 (lcs, role-depth bounded lcs). *The* least common subsumer *of two $\mathcal{L}_s$-concepts C_1, C_2 w.r.t. an $\mathcal{L}_s$-TBox $\mathcal{T}$ (written: $lcs_{\mathcal{T}}(C_1, C_2)$) is the $\mathcal{L}_t$-concept description D s.t.:*

1. $C_1 \sqsubseteq_{\mathcal{T}} D$ and $C_2 \sqsubseteq_{\mathcal{T}} D$, and
2. for all $\mathcal{L}_t$-concepts E, $C_1 \sqsubseteq_{\mathcal{T}} E$ and $C_2 \sqsubseteq_{\mathcal{T}} E$ implies $D \sqsubseteq_{\mathcal{T}} E$.

Let $k \in \mathbb{N}$. If the concept D has a role-depth up to k and Condition 2 holds for all such E with role-depth up to k, then D is the role-depth bounded lcs *(k-$lcs_{\mathcal{T}}(C_1, C_2)$) of C_1 and C_2 w.r.t. $\mathcal{T}$ and k.*

The role-depth bounded lcs is unique up to equivalence, thus we speak of *the k-lcs*. In contrast, common subsumers need not be unique. Note that for target DLs that offer disjunction, the lcs is always trivial: $lcs(C_1, C_2) = C_1 \sqcup C_2$. Thus target DLs without disjunction may yield more informative lcs.

Similarly to the lcs, the msc does not need to exist if computed w.r.t. cyclic ABoxes. Again we compute here approximations of the msc by limiting the role-depth of the resulting concept as suggested in [12].

Definition 2. *Let $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ be a KB written in $\mathcal{L}_s$ and a be an individual from $\mathcal{A}$. An $\mathcal{L}_t$-concept description C is the* most specific concept of a w.r.t. $\mathcal{K}$ *(written $msc_{\mathcal{K}}(a)$) if it satisfies:*

1. $\mathcal{K} \models C(a)$, and
2. *for all $\mathcal{L}_t$-concepts D, $\mathcal{K} \models D(a)$ implies $C \sqsubseteq_{\mathcal{T}} D$.*

If the concept C has a role-depth up to k and Condition 2 holds for all such D with role-depth up to k, then C is the role depth bounded msc *of a w.r.t. $\mathcal{K}$ and k (k-$msc_{\mathcal{K}}(a)$).*

The msc and the *k-msc* are unique up to equivalence in $\mathcal{EL}$ and $\mathcal{ELOR}$. In $\mathcal{ELOR}$ the msc is trivial, since $msc_{\mathcal{K}}(a) = \{a\}$. Thus we consider in this paper a more interesting case, where the target DL $\mathcal{L}_t$ for the resulting concept is a less expressive one without nominals, namely $\mathcal{EL}$ or $\mathcal{ELR}$.

3 The *k*-lcs in $\mathcal{ELOR}$

The algorithms to compute the role-depth bounded lcs are based on completion-based classification algorithms for the corresponding DL. For the DL $\mathcal{ELOR}$, a consequence-based algorithm for classification of TBoxes was presented in [11], building upon the completion algorithm developed in [3]. The completion algorithm presented next adapts the ideas of the complete algorithm.

3.1 Completion Algorithm for $\mathcal{ELOR}$-TBoxes

The completion algorithms work on normalized TBoxes. We define for $\mathcal{ELOR}$ the set of *basic concepts* for a TBox $\mathcal{T}$:

$$\mathsf{BC}_{\mathcal{T}} = (\mathsf{Sig}(\mathcal{T}) \cap (N_C \cup N_I)) \cup \{\top\}.$$

Let $\mathcal{T}$ be an $\mathcal{ELOR}$-TBox and $A, A_1, A_2, B \in \mathsf{BC}_{\mathcal{T}}$; then $\mathcal{T}$ is in *normal form* if

- each GCI in $\mathcal{T}$ is of the form: $A \sqsubseteq B$, $A_1 \sqcap A_2 \sqsubseteq B$, $A \sqsubseteq \exists r.B$, or $\exists r.A \sqsubseteq B$.
- each RIA in $\mathcal{T}$ is of the form: $r \sqsubseteq s$ or $r_1 \circ r_2 \sqsubseteq s$.

Every $\mathcal{ELOR}$-TBox can be transformed into normal form in linear time by applying a set of normalization rules given in [3]. These normalization rules essentially introduce new named concepts for complex concepts used in GCIs or new roles used in complex RIAs.

Before describing the completion algorithm in detail, we introduce the reachability relation $\leadsto_R$, which plays a fundamental role in the correct treatment of nominals in TBox classification algorithms [3, 11].

Definition 3 ($\leadsto_R$). *Let $\mathcal{T}$ be an $\mathcal{ELOR}$-TBox in normal form, $G \in N_C$ a concept name, and $D \in \mathsf{BC}_\mathcal{T}$. $G \leadsto_R D$ iff there exist roles $r_1, \ldots, r_n \in N_R$ and basic concepts $A_0, \ldots, A_n, B_0, \ldots, B_n \in \mathsf{BC}_\mathcal{T}$, $n \geq 0$ such that $A_i \sqsubseteq_\mathcal{T} B_i$ for all $0 \leq i \leq n$, $B_{i-1} \sqsubseteq \exists r_i.A_i \in \mathcal{T}$ for all $1 \leq i \leq n$, A_0 is either G or a nominal, and $B_n = D$.*

Informally, the concept name D is reachable from G if there is a chain of existential restrictions starting from G or a nominal and ending in D. This implies that, for $G \leadsto_R D$, if the interpretation of G is not empty, then the interpretation of D cannot be empty either. This in turn causes additional subsumption relationships to hold. Note that, if D is reachable from a nominal, then $G \leadsto_R D$ holds for all concept names G, since the interpretation of D can never be empty.

The basic idea of completion algorithms in general is to generate canonical models of the TBox. To this end, the elements of the interpretation domain are represented by named concepts or nominals from the normalized TBox. These elements are then related via roles according to the existential restrictions derived for the TBox. More precisely, let $\mathcal{T}$ be a normalized TBox, $G \in \mathsf{Sig}(\mathcal{T}) \cap N_C \cup \{\top\}$ and $A \in \mathsf{BC}_\mathcal{T}$, we introduce a *completion set* $S^G(A)$. We store all basic concepts that subsume a basic concept A in the completion set $S^A(A)$ and all basic concepts B for which $\exists r.B$ subsumes A in the completion set $S^A(A, r)$. These completion sets are then extended using a set of rules. However, the algorithm needs to keep track also of completion sets of the form $S^G(A)$ and $S^G(A, r)$ for every $G \in (\mathsf{Sig}(\mathcal{T}) \cap N_C) \cup \{\top\}$, since the non-emptiness of an interpretation of a concept G may imply additional subsumption relationships for A. The completion set $S^G(A)$ therefore stores all basic concepts that subsume A under the assumption that G is not empty. Similarly $S^G(A, r)$ stores all concepts B for which $\exists r.B$ subsumes A under the same assumption.

For every $G \in (\mathsf{Sig}(\mathcal{T}) \cap N_C) \cup \{\top\}$, every basic concept A and every role name r, the completion sets are initialized as $S^G(A) = \{A, \top\}$ and $S^G(A, r) = \emptyset$. These sets are then extended by applying the completion rules shown in Figure 1 (adapted from [11]) exhaustively.

To compute the reachability relation $\leadsto_R$ used in rule **OR7**, the algorithm can use Definition 3 with all previously derived subsumption relationships; that is, $A_i \sqsubseteq B_i$ if it finds $B_i \in S^{A_i}(A_i)$. Thus the computation of $\leadsto_R$ and the application of the completion rules need to be carried out simultaneously.

It can be shown that the algorithm terminates in polynomial time, and is sound and complete for classifying the TBox $\mathcal{T}$. In particular, when no rules are applicable anymore the completion sets have the following properties.

Proposition 1. *Let $\mathcal{T}$ be an $\mathcal{ELOR}$-TBox in normal form, $C, D \in \mathsf{BC}_\mathcal{T}$, $r \in \mathsf{Sig}(\mathcal{T}) \cap N_R$, and $G = C$ if $C \in N_C$ and $G = \top$ otherwise. Then, the following properties hold:*

$C \sqsubseteq_\mathcal{T} D \quad$ *iff* $\quad D \in S^G(C)$, *and*
$C \sqsubseteq_\mathcal{T} \exists r.D \quad$ *iff* $\quad$ *there exists* $E \in \mathsf{BC}_\mathcal{T}$ *such that* $E \in S^G(C, r)$ *and* $D \in S^G(E)$.

We now show how to use these completion sets for computing the role-depth bounded lcs for $\mathcal{ELOR}$-concept w.r.t. a general $\mathcal{ELOR}$-TBox.

$$
\boxed{
\begin{array}{l}
\textbf{OR1 } \text{If } A_1 \in S^G(A),\ A_1 \sqsubseteq B \in \mathcal{T} \text{ and } B \notin S^G(A), \\
\qquad \text{then } S^G(A) := S^G(A) \cup \{B\} \\[4pt]
\textbf{OR2 } \text{If } A_1, A_2 \in S^G(A),\ A_1 \sqcap A_2 \sqsubseteq B \in \mathcal{T} \text{ and } B \notin S^G(A), \\
\qquad \text{then } S^G(A) := S^G(A) \cup \{B\} \\[4pt]
\textbf{OR3 } \text{If } A_1 \in S^G(A),\ A_1 \sqsubseteq \exists r.B \in \mathcal{T} \text{ and } B \notin S^G(A,r), \\
\qquad \text{then } S^G(A,r) := S^G(A,r) \cup \{B\} \\[4pt]
\textbf{OR4 } \text{If } B \in S^G(A,r),\ B_1 \in S^G(B),\ \exists r.B_1 \sqsubseteq C \in \mathcal{T} \text{ and } C \notin S^G(A), \\
\qquad \text{then } S^G(A) := S^G(A) \cup \{C\} \\[4pt]
\textbf{OR5 } \text{If } B \in S^G(A,r),\ r \sqsubseteq s \in \mathcal{T} \text{ and } B \notin S^G(A,s), \\
\qquad \text{then } S^G(A,s) := S^G(A,s) \cup \{B\} \\[4pt]
\textbf{OR6 } \text{If } B \in S^G(A,r_1),\ C \in S^G(B,r_2),\ r_1 \circ r_2 \sqsubseteq s \in \mathcal{T} \text{ and } C \notin S^G(A,s), \\
\qquad \text{then } S^G(A,s) := S^G(A,s) \cup \{C\} \\[4pt]
\textbf{OR7 } \text{If } \{a\} \in S^G(A_1) \cap S^G(A_2),\ G \leadsto_R A_2, \text{ and } A_2 \notin S^G(A_1), \\
\qquad \text{then } S^G(A_1) := S^G(A_1) \cup \{A_2\}
\end{array}
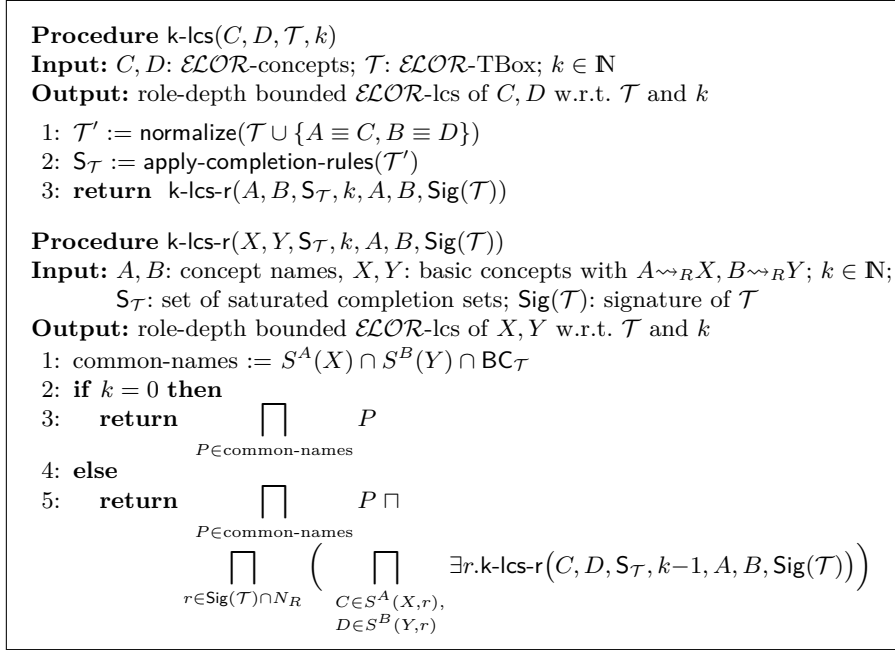}
$$

Fig. 1. Completion rules for $\mathcal{ELOR}$

3.2 Computing the Role-depth Bounded $\mathcal{ELOR}$-lcs

In order to compute the role-depth bounded lcs of two $\mathcal{ELOR}$-concepts C and D, we extend the methods from [16] for $\mathcal{EL}$-concepts and from [9] for $\mathcal{ELR}$-concepts, where we compute the cross-product of the tree unravelings of the canonical model represented by the completion sets for C and D up to the role-depth k. Clearly, in the presence of nominals, the right completion sets need to be chosen that preserve the non-emptiness of the interpretation of concepts derived by $\leadsto_R$.

An algorithm that computes the role-depth bounded $\mathcal{ELOR}$-lcs using completion sets is shown in Figure 2. In the first step, the algorithm introduces two new concept names A and B as abbreviations for the (possibly complex) concepts C and D, and the augmented TBox is normalized. The completion sets are then initialized and the completion rules from Figure 1 are applied exhaustively, yielding the saturated completion sets $\mathsf{S}_\mathcal{T}$. In the recursive procedure k-lcs-r for concepts A and B, we first obtain all the basic concepts that subsume both A and B from the sets $S^A(A)$ and $S^B(B)$. For every role name r, the algorithm then recursively computes the $(k-1)$-lcs of the concepts A' and B' in the subsumer sets $S^A(A,r)$ and $S^B(B,r)$, i.e. for which $A \sqsubseteq_\mathcal{T} \exists r.A'$ and $B \sqsubseteq_\mathcal{T} \exists r.B'$. These concepts are added as existential restrictions to the k-lcs.

The algorithm only introduces concept and role names that occur in the original TBox $\mathcal{T}$. Therefore those names introduced by the normalization are not used in the concept for the k-lcs and an extra denormalization step as in [16, 9] is not necessary.

Notice that for every pair (A', B') of r-successors of A and B it holds that $A \leadsto_R A'$ and $B \leadsto_R B'$. Intuitively, we are assuming that the interpretation of both A and B is not empty. This in turn causes the interpretation of $\exists r.A'$ and $\exists r.B'$ to be not empty, either. Thus, it suffices to consider the completion

```
Procedure k-lcs(C, D, T, k)
Input: C, D: ELOR-concepts; T: ELOR-TBox; k ∈ ℕ
Output: role-depth bounded ELOR-lcs of C, D w.r.t. T and k
 1: T' := normalize(T ∪ {A ≡ C, B ≡ D})
 2: S_T := apply-completion-rules(T')
 3: return  k-lcs-r(A, B, S_T, k, A, B, Sig(T))

Procedure k-lcs-r(X, Y, S_T, k, A, B, Sig(T))
Input: A, B: concept names, X, Y: basic concepts with A⤳_R X, B⤳_R Y; k ∈ ℕ;
         S_T: set of saturated completion sets; Sig(T): signature of T
Output: role-depth bounded ELOR-lcs of X, Y w.r.t. T and k
 1: common-names := S^A(X) ∩ S^B(Y) ∩ BC_T
 2: if k = 0 then
 3:    return    ⊓       P
              P∈common-names
 4: else
 5:    return    ⊓       P ⊓
              P∈common-names

           ⊓       (  ⊓    ∃r.k-lcs-r(C, D, S_T, k−1, A, B, Sig(T)))
       r∈Sig(T)∩N_R    C∈S^A(X,r),
                       D∈S^B(Y,r)
```

Fig. 2. Computation algorithm for role-depth bounded $\mathcal{ELOR}$-lcs.

sets S^A and S^B, without the need to additionally compute $S^{A'}$ and $S^{B'}$, or the completion sets S^C for any other basic concept C encountered during the recursive computation of the k-lcs. This allows for a goal-oriented optimization in cases where there is no need to classify the full TBox.
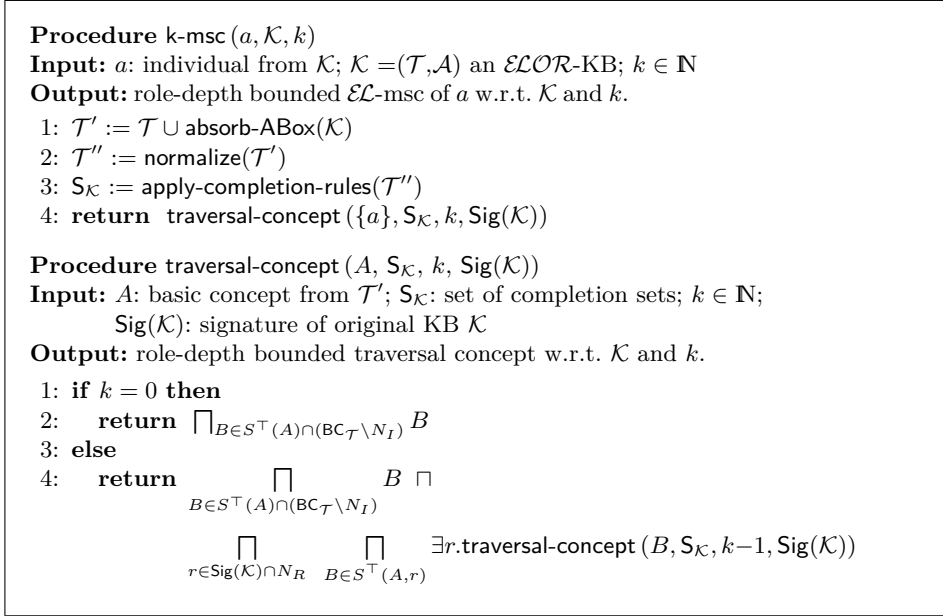
3.3   Computing the Role-depth Bounded msc w.r.t. $\mathcal{ELOR}$-KBs

We now turn our attention to the other generalization inference: the computation of the most specific concept representing a given individual. Recall that, since $\mathcal{ELOR}$ allows the use of nominals, computing the (exact) $\mathcal{ELOR}$-msc for a given individual is a trivial task: the most specific $\mathcal{ELOR}$-concept describing an individual $a \in N_I$ is simply the nominal $\{a\}$. However, it may be of interest to compute the msc w.r.t. a less expressive target DL. Next, we describe how to compute the depth-bounded $\mathcal{EL}$-msc of an individual w.r.t. an $\mathcal{ELOR}$-KB.

As we have defined them, KBs consist of two parts: the TBox, which represents the conceptual knowledge of the domain, and the ABox, which states information about individuals. In the presence of nominals, this division between concepts and individuals is blurred. In fact, it is possible to simulate ABox assertions using GCIs as described by the following proposition.

Lemma 1. *An interpretation $\mathcal{I}$ satisfies the concept assertion $C(a)$ iff it satisfies the GCI $\{a\} \sqsubseteq C$. It satisfies the role assertion $r(a, b)$ iff it satisfies the GCI $\{a\} \sqsubseteq \exists r.\{b\}$.*

Procedure k-msc $(a, \mathcal{K}, k)$
Input: a: individual from $\mathcal{K}$; $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ an $\mathcal{ELOR}$-KB; $k \in \mathbb{N}$
Output: role-depth bounded $\mathcal{EL}$-msc of a w.r.t. $\mathcal{K}$ and k.

1: $\mathcal{T}' := \mathcal{T} \cup \mathsf{absorb\text{-}ABox}(\mathcal{K})$
2: $\mathcal{T}'' := \mathsf{normalize}(\mathcal{T}')$
3: $\mathsf{S}_\mathcal{K} := \mathsf{apply\text{-}completion\text{-}rules}(\mathcal{T}'')$
4: **return** $\mathsf{traversal\text{-}concept}(\{a\}, \mathsf{S}_\mathcal{K}, k, \mathsf{Sig}(\mathcal{K}))$

Procedure traversal-concept $(A, \mathsf{S}_\mathcal{K}, k, \mathsf{Sig}(\mathcal{K}))$
Input: A: basic concept from $\mathcal{T}'$; $\mathsf{S}_\mathcal{K}$: set of completion sets; $k \in \mathbb{N}$;
$\quad\quad$ $\mathsf{Sig}(\mathcal{K})$: signature of original KB $\mathcal{K}$
Output: role-depth bounded traversal concept w.r.t. $\mathcal{K}$ and k.

1: **if** $k = 0$ **then**
2: $\quad$ **return** $\bigsqcap_{B \in S^\top(A) \cap (\mathsf{BC}_\mathcal{T} \setminus N_I)} B$
3: **else**
4: $\quad$ **return** $\displaystyle\bigsqcap_{B \in S^\top(A) \cap (\mathsf{BC}_\mathcal{T} \setminus N_I)} B \; \sqcap$

$\quad\quad\quad\quad \displaystyle\bigsqcap_{r \in \mathsf{Sig}(\mathcal{K}) \cap N_R} \bigsqcap_{B \in S^\top(A, r)} \exists r.\mathsf{traversal\text{-}concept}(B, \mathsf{S}_\mathcal{K}, k-1, \mathsf{Sig}(\mathcal{K}))$

Fig. 3. Computation algorithm for the role-depth bounded $\mathcal{EL}$-msc w.r.t. $\mathcal{ELOR}$-KBs.

Using this result, we can 'absorb' the ABox into the TBox and restrict our attention to reasoning w.r.t. TBoxes only, without losing generality. Figure 3 describes the algorithm for computing the $\mathcal{EL}$-k-msc w.r.t. an $\mathcal{ELOR}$-KB.

As before, correctness of this algorithm is a consequence of the invariants described by Proposition 1. The set $S^\top(\{a\})$ contains all the basic concepts that subsume the nominal $\{a\}$; that is, all concepts whose interpretation must contain the individual $a^\mathcal{I}$. Likewise, $S^\top(\{a\}, r)$ contains all the existential restrictions subsuming $\{a\}$. Thus, a recursive conjunction of all these subsumers provides the most specific representation of the individual a.

Since the target language is $\mathcal{EL}$, no nominals may be included in the output. However, the recursion includes also the $\mathcal{EL}$-msc of the removed nominals, hence indeed providing the most specific $\mathcal{EL}$ representation of the input individual. As in the computation of the lcs presented above, the only completion sets relevant for computing the msc are those of the form $S^\top(A)$ and $S^\top(A, r)$. Once again, this means that it is possible to implement a goal-oriented approach that computes only these sets, as needed, when building the msc for a given individual.

In this section we have shown how to compute generalization inferences with a bounded role-depth in the DL $\mathcal{ELOR}$ that extends $\mathcal{EL}$ by allowing nominals and complex role inclusion axioms in the KB. With the exception of data-types and disjointness axioms, this covers the full expressivity of the OWL 2 EL profile of the standard ontology language OWL 2. Given its status as W3C standard, it is likely that more and bigger ontologies built using this profile, thus the gener-

alization inferences investigated in this paper and their computation algorithms for approximation will become more useful to ontology engineers. In fact, there already exist ontologies that use nominals in their representation. For example, the FMA ontology [18] is written in $\mathcal{ELOR}$ and currently contains 85 nominals.

4   Conclusions

We have studied reasoning services in extensions of the light-weight description logic $\mathcal{EL}$ by nominals and role inclusions, which yields the DL $\mathcal{ELOR}$. One of the characterizing features of $\mathcal{EL}$ and its extension $\mathcal{ELOR}$ is that they allow for polynomial time reasoning. Efficient reasoning becomes expedient when dealing with huge knowledge bases such as the biomedical ontologies SNOMED and the Gene Ontology. Additionally, $\mathcal{ELOR}$ covers a large part of the OWL 2 EL profile. Given its status as a W3C recommendation, it is likely that the usage of the $\mathcal{EL}$-family of DLs becomes more widespread in the future.

Especially for the huge ontologies written in extensions of $\mathcal{EL}$, tools that aid the user with the construction and maintenance of the knowledge base are necessary. As previous work has shown, the generalization inferences lcs and msc can be effectively used for such tasks. Besides this, the generalizations can be used as a basis for other inferences, like the construction of semantic similarity measures and information retrieval procedures.

The contributions of the paper are manyfold. First, we have given a completion algorithm for $\mathcal{ELOR}$ knowledge bases, inspired by a consequence-based classification algorithm for $\mathcal{EL}$ with nominals [11]. This completion algorithm is then employed to extend the algorithms for computing approximations of the lcs and of the msc for the DL $\mathcal{ELOR}$. In general, the lcs and msc do not need to exist, even for $\mathcal{EL}$, thus we approximate them by limiting the role-depth of the resulting concept description, up to a maximal bound specified by the user.

We extended here the computation algorithm of the k-lcs to the DL $\mathcal{ELOR}$, using the new completion algorithm, by allowing nominals as part of the resulting concept. Since the k-msc is trivial in $\mathcal{ELOR}$ due to nominals, we give a computation algorithm for the k-msc for the target language $\mathcal{EL}$, which works for $\mathcal{ELOR}$-KBs. Using these algorithms, the generalization inferences can be used for a large set of ontologies built for the OWL 2 EL profile. Both algorithms have the property that, if the exact lcs or msc exist, then our algorithms compute the exact solution for a sufficiently large role-depth bound k. Such a k can be computed for $\mathcal{EL}$ using the necessary and sufficient conditions for the existence of the lsc and msc given in [21].

As future work we intend to study methods of finding these generalizations in further extensions of $\mathcal{EL}$. Initial steps in this direction have been made by considering $\mathcal{EL}$ with subjective probability constructors [17]. In a different direction, we also intend to implement a system that can compute the lcs and the msc, by modifying and improving existing completion-based reasoners.

References

1. F. Baader. Least common subsumers and most specific concepts in a description logic with existential restrictions and terminological cycles. In G. Gottlob and T. Walsh, editors, *Proc. of the 18th Int. Joint Conf. on Artificial Intelligence (IJCAI-03)*, pages 325–330. Morgan Kaufmann, 2003.

2. F. Baader. Terminological cycles in a description logic with existential restrictions. In G. Gottlob and T. Walsh, editors, *Proc. of the 18th Int. Joint Conf. on Artificial Intelligence (IJCAI-03)*, pages 319–324. Morgan Kaufmann, 2003.

3. F. Baader, S. Brandt, and C. Lutz. Pushing the $\mathcal{EL}$ envelope. In *Proc. of the 19th Int. Joint Conf. on Artificial Intelligence (IJCAI-05)*, Edinburgh, UK, 2005. Morgan-Kaufmann Publishers.

4. F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press, 2003.

5. F. Baader, R. Küsters, and R. Molitor. Computing least common subsumers in description logics with existential restrictions. In T. Dean, editor, *Proc. of the 16th Int. Joint Conf. on Artificial Intelligence (IJCAI-99)*, pages 96–101, Stockholm, Sweden, 1999. Morgan Kaufmann, Los Altos.

6. A. Borgida, T. Walsh, and H. Hirsh. Towards measuring similarity in description logics. In *Proc. of the 2005 Description Logic Workshop (DL 2005)*, volume 147 of *CEUR Workshop*, 2005.

7. S. Brandt and A.-Y. Turhan. Using non-standard inferences in description logics — what does it buy me? In G. Görz, V. Haarslev, C. Lutz, and R. Möller, editors, *Proc. of the 2001 Applications of Description Logic Workshop (ADL 2001)*, number 44 in CEUR Workshop, Vienna, Austria, September 2001. RWTH Aachen. See `http://CEUR-WS.org/Vol-44/`.

8. C. d'Amato, N. Fanizzi, and F. Esposito. A semantic similarity measure for expressive description logics. In *Pro.c of Convegno Italiano di Logica Computazionale, CILC05*, 2005.

9. A. Ecke and A.-Y. Turhan. Role-depth bounded least common subsumers for $\mathcal{EL}^+$ and $\mathcal{ELI}$. In Y. Kazakov, D. Lembo, and F. Wolter, editors, *Proc. of the 2012 Description Logic Workshop (DL 2012)*, volume 846 of *CEUR Workshop Proceedings*. CEUR Workshop, 2012.

10. Y. Kazakov, M. Krötzsch, and F. Simančík. ELK reasoner: Architecture and evaluation. In I. Horrocks, M. Yatskevich, and E. Jimenez-Ruiz, editors, *Proceedings of the OWL Reasoner Evaluation Workshop (ORE'12)*, volume 858 of *CEUR Workshop*. CEUR-WS.org, 2012.

11. Y. Kazakov, M. Krötzsch, and F. Simančík. Practical reasoning with nominals in the $\mathcal{EL}$ family of description logics. In G. Brewka, T. Eiter, and S. A. McIlraith, editors, *Proc. of the 12th Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR-12)*, pages 264–274. AAAI Press, 2012.

12. R. Küsters and R. Molitor. Approximating most specifc concepts in description logics with existential restrictions. In F. Baader, G. Brewka, and T. Eiter, editors, *Proc. of the 24th German Annual Conf. on Artificial Intelligence (KI'01)*, volume 2174 of *Lecture Notes In Artificial Intelligence*, pages 33–47. Springer, 2001.

13. K. Lehmann and A.-Y. Turhan. A framework for semantic-based similarity measures for $\mathcal{ELH}$-concepts. In L. F. del Cerro, A. Herzig, and J. Mengin, editors, *Proceedings of the 13th European Conference on Logics in Artificial Intelligence*, Lecture Notes in Artificial Intelligence, pages 307–319. Springer Verlag, 2012.

14. J. Mendez. jCel: A modular rule-based reasoner. In *In Proc. of the 1st Int. Workshop on OWL Reasoner Evaluation (ORE'12)*, volume 858 of *CEUR Workshop*, 2012.

15. B. Motik, B. Cuenca Grau, I. Horrocks, Z. Wu, A. Fokoue, and C. Lutz. OWL 2 web ontology language profiles. W3C Recommendation, 27 October 2009. `http://www.w3.org/TR/2009/REC-owl2-profiles-20091027/`.

16. R. Peñaloza and A.-Y. Turhan. A practical approach for computing generalization inferences in $\mathcal{EL}$. In M. Grobelnik and E. Simperl, editors, *Proc. of the 8th European Semantic Web Conf. (ESWC'11)*, Lecture Notes in Computer Science. Springer, 2011.

17. R. Peñaloza and A.-Y. Turhan. Instance-based non-standard inferences in $\mathcal{EL}$ with subjective probabilities. In F. Bobillo, P. C. G. Costa, C. d'Amato, N. Fanizzi, K. B. Laskey, K. J. Laskey, T. Lukasiewicz, M. Nickles, and M. Pool, editors, *Uncertainty Reasoning for the Semantic Web II*, number 7123 in Lecture Notes in Computer Science, pages 80–98. Springer, 2013.

18. C. Rosse and J. L. V. Mejino. A reference ontology for biomedical informatics: the foundational model of anatomy. *Journal of Biomedical Informatics*, 36:478–500, 2003.

19. K. Spackman. Managing clinical terminology hierarchies using algorithmic calculation of subsumption: Experience with SNOMED-RT. *Journal of the American Medical Informatics Assoc.*, 2000. Fall Symposium Special Issue.

20. W3C OWL Working Group. OWL 2 web ontology language document overview. W3C Recommendation, 27th October 2009. `http://www.w3.org/TR/2009/REC-owl2-overview-20091027/`.

21. B. Zarrieš and A.-Y. Turhan. Most specific generalizations w.r.t. general $\mathcal{EL}$-tboxes. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI'13)*, Beijing, China, 2013. AAAI Press. To appear.