UNIFICATION IN THE DESCRIPTION LOGIC $\mathcal{EL}$

FRANZ BAADER AND BARBARA MORAWSKA

Theoretical Computer Science, TU Dresden, Germany
e-mail address: {baader,morawska}@tcs.inf.tu-dresden.de

ABSTRACT. The Description Logic $\mathcal{EL}$ has recently drawn considerable attention since, on the one hand, important inference problems such as the subsumption problem are polynomial. On the other hand, $\mathcal{EL}$ is used to define large biomedical ontologies. Unification in Description Logics has been proposed as a novel inference service that can, for example, be used to detect redundancies in ontologies. The main result of this paper is that unification in $\mathcal{EL}$ is decidable. More precisely, $\mathcal{EL}$-unification is NP-complete, and thus has the same complexity as $\mathcal{EL}$-matching. We also show that, w.r.t. the unification type, $\mathcal{EL}$ is less well-behaved: it is of type zero, which in particular implies that there are unification problems that have no finite complete set of unifiers.

1. INTRODUCTION

Description logics (DLs) [6] are a family of logic-based knowledge representation formalisms, which can be used to represent the conceptual knowledge of an application domain in a structured and formally well-understood way. They are employed in various application domains, such as natural language processing, configuration of technical systems, databases, and biomedical ontologies, but their most notable success so far is the adoption of the DL-based language OWL [20] as standard ontology language for the semantic web.

In DLs, concepts are formally described by *concept terms*, i.e., expressions that are built from concept names (unary predicates) and role names (binary predicates) using concept constructors. The expressivity of a particular DL is determined by which concept constructors are available in it. From a semantic point of view, concept names and concept terms represent sets of individuals, whereas roles represent binary relations between individuals. For example, using the concept name Woman, and the role name child, the concept of *women having a daughter* can be represented by the concept term

$$\mathsf{Woman} \sqcap \exists\, \mathsf{child}.\mathsf{Woman},$$

and the concept of *women having only daughters* by

$$\mathsf{Woman} \sqcap \forall\, \mathsf{child}.\mathsf{Woman}.$$

Knowledge representation systems based on DLs provide their users with various inference services that allow them to deduce implicit knowledge from the explicitly represented knowledge. An important inference problem solved by DL systems is the subsumption problem: the subsumption algorithm allows one to determine subconcept-superconcept relationships. For example, the concept term Woman subsumes the concept term Woman $\sqcap\,\exists$ child.Woman since all instances of the second term are also instances of the first term, i.e., the second term is always interpreted as a subset of the first term. With the help of the subsumption algorithm, a newly introduced concept term can automatically be placed at the correct position in the hierarchy of the already existing concept terms.

Two concept terms C, D are *equivalent* ($C \equiv D$) if they subsume each other, i.e., if they always represent the same set of individuals. For example, the terms $\forall$ child.Rich $\sqcap$ $\forall$ child.Woman and $\forall$ child.(Rich $\sqcap$ Woman) are equivalent since the value restriction operator ($\forall r.C$) distributes over the conjunction operator ($\sqcap$). If we replace the value restriction operator by the existential restriction operator ($\exists r.C$), then this equivalence no longer holds. However, for this operator, we still have the equivalence

$$\exists\,\text{child.Rich}\,\sqcap\,\exists\,\text{child.(Woman}\,\sqcap\,\text{Rich}) \equiv \exists\,\text{child.(Woman}\,\sqcap\,\text{Rich}).$$

The equivalence test can, for example, be used to find out whether a concept term representing a particular notion has already been introduced, thus avoiding multiple introduction of the same concept into the concept hierarchy. This inference capability is very important if the knowledge base containing the concept terms is very large, evolves during a long time period, and is extended and maintained by several knowledge engineers. However, testing for equivalence of concepts is not always sufficient to find out whether, for a given concept term, there already exists another concept term in the knowledge base describing the same notion. On the one hand, different knowledge engineers may use different names for concepts, like Male versus Masculine. On the other hand, they may model on different levels of granularity. For example, assume that one knowledge engineer has defined the concept of *men loving fast cars* by the concept term

$$\text{Human}\,\sqcap\,\text{Male}\,\sqcap\,\exists\,\text{loves.Sports\_car}.$$

A second knowledge engineer might represent this notion in a somewhat different way, e.g., by using the concept term

$$\text{Man}\,\sqcap\,\exists\,\text{loves.(Car}\,\sqcap\,\text{Fast}).$$

These two concept terms are not equivalent, but they are meant to represent the same concept. The two terms can obviously be made equivalent by substituting the concept name Sports_car in the first term by the concept term Car $\sqcap$ Fast and the concept name Man in the second term by the concept term Human $\sqcap$ Male. This leads us to *unification of concept terms*, i.e., the question whether two concept terms can be made equivalent by applying an appropriate substitution, where a substitution replaces (some of the) concept names by concept terms. Of course, it is not necessarily the case that unifiable concept terms are meant to represent the same notion. A unifiability test can, however, suggest to the knowledge engineer possible candidate terms. A *unifier* (i.e., a substitution whose application makes the two terms equivalent) then proposes appropriate definitions for the concept names. In our example, we know that, if we define Man as Human $\sqcap$ Male and Sports_car as Car $\sqcap$ Fast, then the concept terms Human $\sqcap$ Male $\sqcap\,\exists$ loves.Sports_car and Man $\sqcap\,\exists$ loves.(Car $\sqcap$ Fast) are equivalent w.r.t. these definitions.

Unification in DLs was first considered in [12] for a DL called $\mathcal{FL}_0$, which has the concept constructors *conjunction* ($\sqcap$), *value restriction* ($\forall\, r.C$), and the *top concept* ($\top$). It was shown that unification in $\mathcal{FL}_0$ is decidable and ExpTime-complete, i.e., given an $\mathcal{FL}_0$-unification problem, we can effectively decide whether it has a solution or not, but in the worst-case, any such decision procedure needs exponential time. This result was extended in [8] to a more expressive DL, which additionally has the role constructor *transitive closure*. Interestingly, the *unification type* of $\mathcal{FL}_0$ had been determined almost a decade earlier in [2]. In fact, as shown in [12], unification in $\mathcal{FL}_0$ corresponds to unification modulo the equational theory of idempotent Abelian monoids with several homomorphisms. In [2] it was shown that, already for a single homomorphism, unification modulo this theory has unification type zero, i.e., there are unification problems for this theory that do not have a minimal complete set of unifiers. In particular, such unification problems cannot have a finite complete set of unifiers.

In this paper, we consider unification in the DL $\mathcal{EL}$. The $\mathcal{EL}$-family consists of inexpressive DLs whose main distinguishing feature is that they provide their users with *existential restrictions* ($\exists\, r.C$) rather than value restrictions ($\forall\, r.C$) as the main concept constructor involving roles. The core language of this family is $\mathcal{EL}$, which has the top concept, conjunction, and existential restrictions as concept constructors. This family has recently drawn considerable attention since, on the one hand, the subsumption problem stays tractable (i.e., decidable in polynomial time) in situations where $\mathcal{FL}_0$, the corresponding DL with value restrictions, becomes intractable: subsumption between concept terms is tractable for both $\mathcal{FL}_0$ and $\mathcal{EL}$ [25, 10], but allowing the use of concept definitions or even more expressive terminological formalisms makes $\mathcal{FL}_0$ intractable [26, 3, 23, 5], whereas it leaves $\mathcal{EL}$ tractable [4, 17, 5]. On the other hand, although of limited expressive power, $\mathcal{EL}$ is nevertheless used in applications, e.g., to define biomedical ontologies. For example, both the large medical ontology SNOMED CT[1] and the Gene Ontology[2] can be expressed in $\mathcal{EL}$, and the same is true for large parts of the medical ontology GALEN [27]. The importance of $\mathcal{EL}$ can also be seen from the fact that the new OWL 2 standard[3] contains a sub-profile OWL 2 EL, which is based on (an extension of) $\mathcal{EL}$.

Unification in $\mathcal{EL}$ has, to the best of our knowledge, not been investigated before, but matching (where one side of the equation(s) to be solved does not contain variables) has been considered in [7, 24]. In particular, it was shown in [24] that the decision problem, i.e., the problem of deciding whether a given $\mathcal{EL}$-matching problem has a matcher or not, is NP-complete. Interestingly, $\mathcal{FL}_0$ behaves better w.r.t. matching than $\mathcal{EL}$: for $\mathcal{FL}_0$, the decision problem is tractable [9]. In this paper, we show that, w.r.t. the unification type, $\mathcal{FL}_0$ and $\mathcal{EL}$ behave the same: just as $\mathcal{FL}_0$, the DL $\mathcal{EL}$ has unification type zero. However, w.r.t. the decision problem, $\mathcal{EL}$ behaves much better than $\mathcal{FL}_0$: $\mathcal{EL}$-unification is NP-complete, and thus has the same complexity as $\mathcal{EL}$-matching.

Regarding unification in DLs that are more expressive than $\mathcal{EL}$ and $\mathcal{FL}_0$, one must look at the literature on unification in modal logics. It is well-known that there is a close connection between modal logics and DLs [6]. For example, the DL $\mathcal{ALC}$, which can be obtained by adding negation to $\mathcal{EL}$ or $\mathcal{FL}_0$, corresponds to the basic (multi-)modal logic K. Decidability of unification in K is a long-standing open problem. Recently, undecidability of unification in some extensions of K (for example, by the universal modality) was shown

---

[1]http://www.ihtsdo.org/snomed-ct/

[2]http://www.geneontology.org/

[3]See http://www.w3.org/TR/owl2-profiles/

in [29]. The undecidability results in [29] also imply undecidability of unification in some expressive DLs (e.g., $\mathcal{SHIQ}$ [21]). The unification types of some modal (and related) logics have been determined by Ghilardi; for example in [19] he shows that K4 and S4 have unification type finitary. Unification in sub-Boolean modal logics (i.e., modal logics that are not closed under all Boolean operations, such as the modal logics corresponding to $\mathcal{EL}$ and $\mathcal{FL}_0$) has, to the best of our knowledge, not been considered in the modal logic literature.

In addition to unification of concept terms as introduced until now, we will also consider unification w.r.t. a so-called acyclic TBox in this article. Until now, we have only talked about concept terms, i.e., complex descriptions of concepts that are built from concept and role names using the concept constructors of the given DL. In applications of DLs, it is, of course, inconvenient to always use such complex descriptions when referring to concepts. For this reason, DLs are usually also equipped with a terminological formalism. In its simplest form, this formalism allows to introduce abbreviations for concept terms. For example, the two concept definitions

$$\mathsf{Mother} \equiv \mathsf{Woman} \sqcap \exists\,\mathsf{child}.\mathsf{Human} \quad \text{and} \quad \mathsf{Woman} \equiv \mathsf{Human} \sqcap \mathsf{Female}$$

introduce the abbreviation Woman for the concept term $\mathsf{Human} \sqcap \mathsf{Female}$ and the abbreviation Mother for the concept term $\mathsf{Human} \sqcap \mathsf{Female} \sqcap \exists\,\mathsf{child}.\mathsf{Human}$. A finite set of such concept definitions is called an *acyclic* TBox if it is unambiguous (i.e., every concept name occurs at most once as left-hand side) and acyclic (i.e., there are no cyclic dependencies between concept definitions). These restrictions ensure that every defined concept (i.e., concept name occurring on the left-hand side of a definition) has a unique expansion to a concept term that it abbreviates. Inference problems like subsumption and unification can also be considered w.r.t. such acyclic TBoxes. As mentioned above, the complexity of the subsumption problem increases for the DL $\mathcal{FL}_0$ if acyclic TBoxes are taken into account [26]. In contrast, for $\mathcal{EL}$, the complexity of the subsumption problem stays polynomial in the presence of acyclic TBoxes. We show that, for unification in $\mathcal{EL}$, adding acyclic TBoxes is also harmless, i.e., unification in $\mathcal{EL}$ *w.r.t. acyclic TBoxes* is also NP-complete.

This article is structured as follows. In the next section, we define the DL $\mathcal{EL}$ and unification in $\mathcal{EL}$ more formally. In Section 3, we recall the characterization of subsumption and equivalence in $\mathcal{EL}$ from [24], and in Section 4 we use this to show that unification in $\mathcal{EL}$ has type zero. In Section 5, we show that unification in $\mathcal{EL}$ is NP-complete. The unification algorithm establishing the complexity upper bound is a typical "guess and then test" NP-algorithm, and thus it is unlikely that a direct implementation of this algorithm will perform well in practice. In Section 6, we introduce a more goal-oriented unification algorithm for $\mathcal{EL}$, in which non-deterministic decisions are only made if they are triggered by "unsolved parts" of the unification problem. In Section 7, we point out that our results for $\mathcal{EL}$-unification imply that unification modulo the equational theory of semilattices with monotone operators [28] is NP-complete and of unification type zero.

More information about Description Logics can be found in [6], and about unification theory in [16]. This article is an extended version of a paper [11] published in the proceedings of the 20th international Conference on Rewriting Techniques and applications (RTA'09). In addition to giving more detailed proofs, we have added the goal-oriented unification algorithm (Section 6) and the treatment of unification modulo acyclic TBoxes (Subsection 2.3).

Name	Syntax	Semantics
concept name	A	$A^{\mathcal{I}} \subseteq \mathcal{D}_{\mathcal{I}}$
role name	r	$r^{\mathcal{I}} \subseteq \mathcal{D}_{\mathcal{I}} \times \mathcal{D}_{\mathcal{I}}$
top-concept	$\top$	$\top^{\mathcal{I}} = \mathcal{D}_{\mathcal{I}}$
conjunction	$C \sqcap D$	$(C \sqcap D)^{\mathcal{I}} = C^{\mathcal{I}} \cap D^{\mathcal{I}}$
existential restriction	$\exists r.C$	$(\exists r.C)^{\mathcal{I}} = \{x \mid \exists y : (x,y) \in r^{\mathcal{I}} \wedge y \in C^{\mathcal{I}}\}$
subsumption	$C \sqsubseteq D$	$C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$
equivalence	$C \equiv D$	$C^{\mathcal{I}} = D^{\mathcal{I}}$

Table 1: Syntax and semantics of $\mathcal{EL}$

2. Unification in $\mathcal{EL}$

In this section, we first define the syntax and semantics of $\mathcal{EL}$-concept terms as well as the subsumption and the equivalence relation on these terms. Then, we introduce unification of $\mathcal{EL}$-concept terms, and finally extend this notion to unification modulo an acyclic TBox.

2.1. **The Description Logic $\mathcal{EL}$.** Starting with a set N_{con} of concept names and a set N_{role} of role names, $\mathcal{EL}$-*concept terms* are built using the following concept constructors: the nullary constructor *top-concept* ($\top$), the binary constructor *conjunction* ($C \sqcap D$), and for every role name $r \in N_{role}$, the unary constructor *existential restriction* ($\exists r.C$). The semantics of $\mathcal{EL}$ is defined in the usual way, using the notion of an interpretation $\mathcal{I} = (\mathcal{D}_{\mathcal{I}}, \cdot^{\mathcal{I}})$, which consists of a nonempty domain $\mathcal{D}_{\mathcal{I}}$ and an interpretation function $\cdot^{\mathcal{I}}$ that assigns binary relations on $\mathcal{D}_{\mathcal{I}}$ to role names and subsets of $\mathcal{D}_{\mathcal{I}}$ to concept terms, as shown in the semantics column of Table 1.

The concept term C *is subsumed by* the concept term D (written $C \sqsubseteq D$) iff $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$ holds for all interpretations $\mathcal{I}$. We say that C *is equivalent to* D (written $C \equiv D$) iff $C \sqsubseteq D$ and $D \sqsubseteq C$, i.e., iff $C^{\mathcal{I}} = D^{\mathcal{I}}$ holds for all interpretations $\mathcal{I}$. The concept term C *is strictly subsumed by* the concept term D (written $C \sqsubset D$) iff $C \sqsubseteq D$ and $C \not\equiv D$. It is well-known that subsumption (and thus also equivalence) of $\mathcal{EL}$-concept terms can be decided in polynomial time [10].

2.2. **Unification of concept terms.** In order to define unification of concept terms, we first introduce the notion of a substitution operating on concept terms. To this purpose, we partition the set of concepts names into a set N_v of concept variables (which may be replaced by substitutions) and a set N_c of concept constants (which must not be replaced by substitutions). Intuitively, N_v are the concept names that have possibly been given another name or been specified in more detail in another concept term describing the same notion. The elements of N_c are the ones of which it is assumed that the same name is used by all knowledge engineers (e.g., standardized names in a certain domain).

A *substitution* σ is a mapping from N_v into the set of all $\mathcal{EL}$-concept terms. This mapping is extended to concept terms in the obvious way, i.e.,

- $\sigma(A) := A$ for all $A \in N_c$,
- $\sigma(\top) := \top$,
- $\sigma(C \sqcap D) := \sigma(C) \sqcap \sigma(D)$, and

- $\sigma(\exists\, r.C) := \exists\, r.\sigma(C)$.

Definition 2.1. An $\mathcal{EL}$-*unification problem* is of the form $\Gamma = \{C_1 \equiv^? D_1, \dots, C_n \equiv^? D_n\}$, where $C_1, D_1, \dots, C_n, D_n$ are $\mathcal{EL}$-concept terms. The substitution σ is a *unifier* (or *solution*) of Γ iff $\sigma(C_i) \equiv \sigma(D_i)$ for $i = 1, \dots, n$. In this case, Γ is called *solvable* or *unifiable*.

When we say that $\mathcal{EL}$-unification is *decidable*, then we mean that the following decision problem is decidable: given an $\mathcal{EL}$-unification problem Γ, decide whether Γ is solvable or not. Accordingly, we say that $\mathcal{EL}$-unification is *NP-complete* if this decision problem is NP-complete.

In the following, we introduce some standard notions from unification theory [16], but formulated for the special case of $\mathcal{EL}$-unification rather than for an arbitrary equational theory. Unifiers can be compared using the instantiation preorder $\preceq$. Let Γ be an $\mathcal{EL}$-unification problem, V the set of variables occurring in Γ, and σ, θ two unifiers of this problem. We define

$$\sigma \preceq \theta \quad \text{iff} \quad \text{there is a substitution } \lambda \text{ such that } \theta(X) \equiv \lambda(\sigma(X)) \text{ for all } X \in V.$$

If $\sigma \preceq \theta$, then we say that θ is an *instance* of σ.

Definition 2.2. Let Γ be an $\mathcal{EL}$-unification problem. The set of substitutions M is called a *complete set of unifiers* for Γ iff it satisfies the following two properties:

(1) every element of M is a unifier of Γ;
(2) if θ is a unifier of Γ, then there exists a unifier $\sigma \in M$ such that $\sigma \preceq \theta$.

The set M is called a *minimal complete set of unifiers* for Γ iff it additionally satisfies

(3) if $\sigma, \theta \in M$, then $\sigma \preceq \theta$ implies $\sigma = \theta$.

The unification type of a given unification problem is determined by the existence and cardinality[4] of such a minimal complete set.

Definition 2.3. Let Γ be an $\mathcal{EL}$-unification problem. This problem has type
- *unitary* iff it has a minimal complete set of unifiers of cardinality 1;
- *finitary* iff it has a finite minimal complete set of unifiers;
- *infinitary* iff it has an infinite minimal complete set of unifiers;
- *zero* iff it does not have a minimal complete set of unifiers.

Note that the set of all unifiers of a given $\mathcal{EL}$-unification problem is always a complete set of unifiers. However, this set is usually infinite and redundant (in the sense that some unifiers are instances of others). For a unitary or finitary $\mathcal{EL}$-unification problem, all unifiers can be represented by a finite complete set of unifiers, whereas for problems of type infinitary or zero this is no longer possible. In fact, if a problem has a finite complete set of unifiers M, then it also has a finite *minimal* complete set of unifiers, which can be obtained by iteratively removing redundant elements from M. For an infinite complete set of unifiers, this approach of removing redundant unifiers may be infinite, and the set reached in the limit need no longer be complete. This is what happens for problems of type zero. The difference between infinitary and type zero is that a unification problem of type zero cannot even have a non-redundant complete set of unifiers, i.e., every complete set of unifiers must contain different unifiers σ, θ such that $\sigma \preceq \theta$. More information on unification type zero can be found in [1].

---

[4]It is easy to see that the cardinality of a minimal complete set of unifiers is uniquely determined by the unification problem.

When we say that $\mathcal{EL}$ *has unification type zero*, we mean that there exists an $\mathcal{EL}$-unification problem that has type zero. Before we can prove in Section 4 that this is indeed the case, we must have a closer look at equivalence in $\mathcal{EL}$ in Section 3. But first, we consider unification modulo acyclic TBoxes.

2.3. **Unification modulo acyclic TBoxes.** A *concept definition* is of the form $A \doteq C$ where A is a concept name and C is a concept term. A *TBox* $\mathcal{T}$ is a finite set of concept definitions such that no concept name occurs more than once on the left-hand side of a concept definition in $\mathcal{T}$. The TBox $\mathcal{T}$ is called *acyclic* if there are no cyclic dependencies between its concept definitions. To be more precise, we say that the concept name A *directly depends on* the concept name B in a TBox $\mathcal{T}$ if $\mathcal{T}$ contains a concept definition $A \doteq C$ and B occurs in C. Let *depends on* be the transitive closure of the relation *directly depends on*. Then $\mathcal{T}$ contains a *terminological cycle* if there is a concept name A that depends on itself. Otherwise, $\mathcal{T}$ is called *acyclic*. Given a TBox $\mathcal{T}$, we call a concept name A a *defined concept* if it occurs as the left-side of a concept definition $A \doteq C$ in $\mathcal{T}$. All other concept names are called *primitive concepts*.

The interpretation $\mathcal{I}$ is a model of the TBox $\mathcal{T}$ iff $A^{\mathcal{I}} = C^{\mathcal{I}}$ holds for all concept definitions $A \doteq C$ in $\mathcal{T}$. Subsumption and equivalence w.r.t. a TBox are defined as follows: $C \sqsubseteq_{\mathcal{T}} D$ ($C \equiv_{\mathcal{T}} D$) iff $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$ ($C^{\mathcal{I}} = D^{\mathcal{I}}$) holds for all models $\mathcal{I}$ of $\mathcal{T}$.

Subsumption and equivalence w.r.t. an acyclic TBox can be reduced to subsumption and equivalence of concept terms (without TBox) by *expanding* the concept terms w.r.t. the TBox: given a concept term C, its expansion $C^{\mathcal{T}}$ w.r.t. the acyclic TBox $\mathcal{T}$ is obtained by exhaustively replacing all defined concept names A occurring on the left-hand side of concept definitions $A \doteq C$ in $\mathcal{T}$ by their defining concept terms C. Given concept terms C, D, we have $C \sqsubseteq_{\mathcal{T}} D$ iff $C^{\mathcal{T}} \sqsubseteq D^{\mathcal{T}}$ [14]. The same is true for equivalence, i.e., $C \equiv_{\mathcal{T}} D$ iff $C^{\mathcal{T}} \equiv D^{\mathcal{T}}$. This expansion process may, however, result in an exponential blow-up [26, 14], and thus this reduction of subsumption and equivalence w.r.t. an acyclic TBox to subsumption and equivalence without a TBox is not polynomial. Nevertheless, in $\mathcal{EL}$, subsumption (and thus also equivalence) w.r.t. acyclic TBoxes can be decided in polynomial time [4].

In our definition of unification modulo acyclic TBoxes, we assume that all defined concepts are concept constants. In fact, defined concepts already have a definition in the given TBox, and thus it does not make sense to introduce new ones for them by unification. In this setting, a *substitution* σ is a mapping from N_v into the set of all $\mathcal{EL}$-concept terms *not containing any defined concepts*.[5] The extension of σ to concept terms is defined as in the previous subsection, and its application to $\mathcal{T}$ is defined as

$$\sigma(\mathcal{T}) := \{A \doteq \sigma(C) \mid A \doteq C \in \mathcal{T}\}.$$

Definition 2.4. An $\mathcal{EL}$-*unification problem modulo an acyclic TBox* is of the form $\Gamma = \{C_1 \equiv^?_{\mathcal{T}} D_1, \ldots, C_n \equiv^?_{\mathcal{T}} D_n\}$, where $C_1, D_1, \ldots, C_n, D_n$ are $\mathcal{EL}$-concept terms, and $\mathcal{T}$ is an acyclic $\mathcal{EL}$-TBox. The substitution σ is a *unifier* (or *solution*) of Γ *modulo* $\mathcal{T}$ iff $\sigma(C_i) \equiv_{\sigma(\mathcal{T})} \sigma(D_i)$ for $i = 1, \ldots, n$. In this case, Γ is called *solvable modulo* $\mathcal{T}$ or *unifiable modulo* $\mathcal{T}$.

Coming back to our example from the introduction, assume that one knowledge engineer has written the concept definition

$$\mathsf{Real\_man} \doteq \mathsf{Human} \sqcap \mathsf{Male} \sqcap \exists\,\mathsf{loves.Sports\_car}.$$

---

[5]This restriction prevents the unifier from introducing cycles into the TBox.

to the TBox, whereas a second one has written the definition

$$\mathsf{Stupid\_man} \doteq \mathsf{Man} \sqcap \exists\,\mathsf{loves}.(\mathsf{Car} \sqcap \mathsf{Fast}),$$

where all the concept names occurring on the left-hand side of these definitions are primitive concepts. Then the substitution that replaces $\mathsf{Sports\_car}$ by $\mathsf{Car} \sqcap \mathsf{Fast}$ and Man by $\mathsf{Human} \sqcap \mathsf{Male}$ is a unifier of $\{\mathsf{Real\_man} \equiv^?_{\mathcal{T}} \mathsf{Stupid\_man}\}$ w.r.t. the TBox $\mathcal{T}$ consisting of these two definitions.

Using expansion, we can reduce unification modulo an acyclic TBox to unification without a TBox. In fact, the following lemma is an easy consequence of the fact that $\sigma(C^{\mathcal{T}}) = \sigma(C)^{\sigma(\mathcal{T})}$ holds for all $\mathcal{EL}$-concept terms C.

Lemma 2.5. *The substitution σ is a unifier of $\{C_1 \equiv^?_{\mathcal{T}} D_1, \ldots, C_n \equiv^?_{\mathcal{T}} D_n\}$ modulo $\mathcal{T}$ iff it is a unifier of $\{C_1^{\mathcal{T}} \equiv^? D_1^{\mathcal{T}}, \ldots, C_n^{\mathcal{T}} \equiv^? D_n^{\mathcal{T}}\}$.*

Since expansion can cause an exponential blow-up, this is not a polynomial reduction. In the remainder of this subsection, we show that there actually exists a polynomial-time reduction of unification modulo an acyclic TBox to unification without a TBox.

We say that the $\mathcal{EL}$-unification problem Γ is in *dag-solved form* if it can be written as $\Gamma = \{X_1 \equiv^? C_1, \ldots, X_n \equiv^? C_n\}$, where $X_1, \ldots, X_n$ are distinct concept variables such that, for all $i \le n$, X_i does not occur in $C_i, \ldots, C_n$. For $i = 1, \ldots, n$, let σ_i be the substitution that maps X_i to C_i and leaves all other variables unchanged. We define the substitution σ_Γ as

$$\sigma_\Gamma(X_i) := \sigma_n(\cdots(\sigma_i(X_i))\cdots)$$

for $i = 1, \ldots, n$, and $\sigma_\Gamma(X) := X$ for all other variables X. The following is an instance of a well-known fact from unification theory [22].

Lemma 2.6. *Let $\Gamma = \{X_1 \equiv^? C_1, \ldots, X_n \equiv^? C_n\}$ be an $\mathcal{EL}$-unification problem in dag-solved form. Then, the set $\{\sigma_\Gamma\}$ is a complete set of unifiers for Γ.*

There is a close relationship between acyclic TBoxes and unification problems in dag-solved form. In fact, if $\mathcal{T}$ is an acyclic TBox, then there is an enumeration $A_1, \ldots, A_n$ of the defined concepts in $\mathcal{T}$ such that $\mathcal{T} = \{A_1 \doteq C_1, \ldots, A_n \doteq C_n\}$ and A_i does not occur in $C_i, \ldots, C_n$. Consequently, the corresponding unification problem

$$\Gamma(\mathcal{T}) := \{A_1 \equiv^? C_1, \ldots, A_n \equiv^? C_n\}$$

(where $A_1, \ldots, A_n$ are now viewed as concept variables) is in dag-solved form. In addition, it is easy to see that, for any $\mathcal{EL}$-concept term C, we have $C^{\mathcal{T}} = \sigma_{\Gamma(\mathcal{T})}(C)$.

Lemma 2.7. *The $\mathcal{EL}$-unification problem $\Gamma = \{C_1 \equiv^?_{\mathcal{T}} D_1, \ldots, C_n \equiv^?_{\mathcal{T}} D_n\}$ is solvable modulo the acyclic TBox $\mathcal{T}$ iff $\{C_1 \equiv^? D_1, \ldots, C_n \equiv^? D_n\} \cup \Gamma(\mathcal{T})$ is solvable.*[6]

Proof. Assume that θ is a unifier of $\Gamma = \{C_1 \equiv^?_{\mathcal{T}} D_1, \ldots, C_n \equiv^?_{\mathcal{T}} D_n\}$ modulo $\mathcal{T}$. Then it is a unifier of $\widehat{\Gamma} := \{C_1^{\mathcal{T}} \equiv^? D_1^{\mathcal{T}}, \ldots, C_n^{\mathcal{T}} \equiv^? D_n^{\mathcal{T}}\}$, by Lemma 2.5. Since $C_i^{\mathcal{T}} = \sigma_{\Gamma(\mathcal{T})}(C_i)$ and $D_i^{\mathcal{T}} = \sigma_{\Gamma(\mathcal{T})}(D_i)$, we have $\widehat{\Gamma} = \{\sigma_{\Gamma(\mathcal{T})}(C_i) \equiv^? \sigma_{\Gamma(\mathcal{T})}(D_i) \mid 1 \le i \le n\}$. Consequently, if we define the substitution τ by setting $\tau(X) := \theta(\sigma_{\Gamma(\mathcal{T})}(X))$ for all concept variables and defined concepts X, then τ is a unifier of $\{C_1 \equiv^? D_1, \ldots, C_n \equiv^? D_n\}$. In addition, since $\sigma_{\Gamma(\mathcal{T})}$ is a unifier of $\Gamma(\mathcal{T})$, τ is also a unifier of $\{C_1 \equiv^? D_1, \ldots, C_n \equiv^? D_n\} \cup \Gamma(\mathcal{T})$.

_____

[6]Note that the defined concepts of $\mathcal{T}$ are treated as concept constants in Γ, and as concept variables in $\{C_1 \equiv^? D_1, \ldots, C_n \equiv^? D_n\} \cup \Gamma(\mathcal{T})$.

Conversely, assume that τ is a unifier of $\{C_1 \equiv^? D_1, \ldots, C_n \equiv^? D_n\} \cup \Gamma(\mathcal{T})$. In particular, this implies that τ is a unifier of $\Gamma(\mathcal{T})$. By Lemma 2.6, $\{\sigma_{\Gamma(\mathcal{T})}\}$ is a complete set of unifiers for $\Gamma(\mathcal{T})$, and thus there is a substitution θ such that $\tau(X) = \theta(\sigma_{\Gamma(\mathcal{T})}(X))$ for all concept variables occurring in the unification problem $\{C_1 \equiv^? D_1, \ldots, C_n \equiv^? D_n\} \cup \Gamma(\mathcal{T})$. Since $C_i^{\mathcal{T}} = \sigma_{\Gamma(\mathcal{T})}(C_i)$ and $D_i^{\mathcal{T}} = \sigma_{\Gamma(\mathcal{T})}(D_i)$, this implies that θ is a unifier of $\{C_1^{\mathcal{T}} \equiv^? D_1^{\mathcal{T}}, \ldots, C_n^{\mathcal{T}} \equiv^? D_n^{\mathcal{T}}\}$, and thus of $\Gamma = \{C_1 \equiv^?_{\mathcal{T}} D_1, \ldots, C_n \equiv^?_{\mathcal{T}} D_n\}$ modulo $\mathcal{T}$, by Lemma 2.5. $\square$

Since the size of $\Gamma(\mathcal{T})$ is basically the same as the size of $\mathcal{T}$, the size of $\Gamma \cup \Gamma(\mathcal{T})$ is linear in the size of Γ and $\mathcal{T}$. Thus, the above lemma provides us with a polynomial-time reduction of $\mathcal{EL}$-unification w.r.t. acyclic TBoxes to $\mathcal{EL}$-unification.

Theorem 2.8. *$\mathcal{EL}$-unification w.r.t. acyclic TBoxes can be reduced in polynomial time to $\mathcal{EL}$-unification.*

3. EQUIVALENCE AND SUBSUMPTION IN $\mathcal{EL}$

In order to characterize equivalence of $\mathcal{EL}$-concept terms, the notion of a reduced $\mathcal{EL}$-concept term is introduced in [24]. A given $\mathcal{EL}$-concept term can be transformed into an equivalent reduced term by applying the following rules modulo associativity and commutativity of conjunction:

$$C \sqcap \top \to C \qquad \text{for all } \mathcal{EL}\text{-concept terms } C$$
$$A \sqcap A \to A \qquad \text{for all concept names } A \in N_{con}$$
$$\exists r.C \sqcap \exists r.D \to \exists r.C \quad \text{for all } \mathcal{EL}\text{-concept terms } C, D \text{ with } C \sqsubseteq D$$

Obviously, these rules are equivalence preserving. We say that the $\mathcal{EL}$-concept term D *is reduced* if none of the above rules is applicable to it (modulo associativity and commutativity of $\sqcap$), and that C *can be reduced to* D if D can be obtained from C by applying the above rules (modulo associativity and commutativity of $\sqcap$). The $\mathcal{EL}$-concept term D is a *reduced form* of C if C can be reduced to D and D is reduced. The following theorem is an easy consequence of Theorem 6.3.1 on page 181 of [24].

Theorem 3.1. *Let C, D be $\mathcal{EL}$-concept terms, and $\widehat{C}, \widehat{D}$ reduced forms of C, D, respectively. Then $C \equiv D$ iff $\widehat{C}$ is identical to $\widehat{D}$ up to associativity and commutativity of $\sqcap$.*

This theorem can also be used to derive a recursive characterization of subsumption in $\mathcal{EL}$. In fact, if $C \sqsubseteq D$, then $C \sqcap D \equiv C$, and thus C and $C \sqcap D$ have the same reduced form. Thus, during reduction, all concept names and existential restrictions of D must be "eaten up" by corresponding concept names and existential restrictions of C.

Corollary 3.2. *Let $C = A_1 \sqcap \ldots \sqcap A_k \sqcap \exists r_1.C_1 \sqcap \ldots \sqcap \exists r_m.C_m$ and $D = B_1 \sqcap \ldots \sqcap B_\ell \sqcap \exists s_1.D_1 \sqcap \ldots \sqcap \exists s_n.D_n$, where $A_1, \ldots, A_k, B_1, \ldots, B_\ell$ are concept names. Then $C \sqsubseteq D$ iff $\{B_1, \ldots, B_\ell\} \subseteq \{A_1, \ldots, A_k\}$ and for every $j, 1 \le j \le n$, there exists an $i, 1 \le i \le m$, such that $r_i = s_j$ and $C_i \sqsubseteq D_j$.*

Note that this corollary also covers the cases where some of the numbers k, ℓ, m, n are zero. The empty conjunction should then be read as $\top$. The following lemma, which is an immediate consequence of this corollary, will be used in our proof that $\mathcal{EL}$ has unification type zero.

Lemma 3.3. *If C, D are reduced $\mathcal{EL}$-concept terms such that $\exists r.D \sqsubseteq C$, then C is either $\top$, or of the form $C = \exists r.C_1 \sqcap \ldots \sqcap \exists r.C_n$ where $n \geq 1$; $C_1, \ldots, C_n$ are reduced and pairwise incomparable w.r.t. subsumption; and $D \sqsubseteq C_1, \ldots, D \sqsubseteq C_n$. Conversely, if C, D are $\mathcal{EL}$-concept terms such that $C = \exists r.C_1 \sqcap \ldots \sqcap \exists r.C_n$ and $D \sqsubseteq C_1, \ldots, D \sqsubseteq C_n$, then $\exists r.D \sqsubseteq C$.*

The following lemma states several other obvious consequences of Corollary 3.2.

Lemma 3.4.
(1) *The existential restriction $\exists r.C$ is reduced iff C is reduced.*
(2) *Let $C_1 \sqcap \ldots \sqcap C_n$ be concept names or existential restrictions. Then the conjunction $C_1 \sqcap \ldots \sqcap C_n$ is reduced iff $C_1, \ldots, C_n$ are reduced and pairwise incomparable w.r.t. subsumption.*
(3) *Let $C = C_1 \sqcap \ldots \sqcap C_m$ and $D = D_1 \sqcap \ldots \sqcap D_n$ be conjunctions of $\mathcal{EL}$-concept terms. If, for all $i, 1 \leq i \leq n$, there exists $j, 1 \leq j \leq m$, such that $C_j \sqsubseteq D_i$, then $C \sqsubseteq D$. If $D_1, \ldots, D_n$ are concept names or existential restrictions, then the implication in the other direction also holds.*

In the proof of decidability of $\mathcal{EL}$-unification, we will make use of the fact that the inverse strict subsumption order is well-founded.

Proposition 3.5. *There is no infinite sequence $C_0, C_1, C_2, C_3, \ldots$ of $\mathcal{EL}$-concept terms such that $C_0 \sqsubset C_1 \sqsubset C_2 \sqsubset C_3 \sqsubset \cdots$.*

Proof. We define the *role depth* of an $\mathcal{EL}$-concept term C as the maximal nesting of existential restrictions in C. Let n_0 be the role depth of C_0. Since $C_0 \sqsubseteq C_i$ for $i \geq 1$, it is an easy consequence of Corollary 3.2 that the role depth of C_i is bounded by n_0, and that C_i contains only concept and role names occurring in C_0. In addition, it is known that, for a given natural number n_0 and finite sets of concept names N_{con} and role names N_{role}, there are, up to equivalence, only finitely many $\mathcal{EL}$-concept terms built using concept names from $\mathcal{C}$ and role names from $\mathcal{R}$ and of a role depth bounded by n_0 [15]. Consequently, there are indices $i < j$ such that $C_i \equiv C_j$. This contradicts our assumption that $C_i \sqsubset C_j$. $\square$

4. An $\mathcal{EL}$-unification problem of type zero

To show that $\mathcal{EL}$ has unification type zero, we exhibit an $\mathcal{EL}$-unification problem that has this type.

Theorem 4.1. *Let X, Y be variables. The $\mathcal{EL}$-unification problem $\Gamma := \{X \sqcap \exists r.Y \equiv^? \exists r.Y\}$ has unification type zero.*

Proof. It is enough to show that any complete set of unifiers for this problem is redundant, i.e., contains two different unifiers that are comparable w.r.t. the instantiation preorder. Thus, let M be a complete set of unifiers for Γ.

First, note that M must contain a unifier that maps X to an $\mathcal{EL}$-concept term not equivalent to $\top$ or $\exists r.\top$. In fact, consider a substitution τ such that $\tau(X) = \exists r.A$ and $\tau(Y) = A$. Obviously, τ is a unifier of Γ. Thus, M must contain a unifier σ such that $\sigma \preceq \tau$. In particular, this means that there is a substitution λ such that $\exists r.A = \tau(X) \equiv \lambda(\sigma(X))$. Obviously, $\sigma(X) \equiv \top$ would imply $\lambda(\sigma(X)) \equiv \top$, and thus $\exists r.A \equiv \top$, which is, however,

not the case. Similarly, $\sigma(X) \equiv \exists r.\top$ would imply $\lambda(\sigma(X)) \equiv \exists r.\top$, and thus $\exists r.A \equiv \exists r.\top$, which is also not the case.

Thus, let $\sigma \in M$ be such that $\sigma(X) \not\equiv \top$ and $\sigma(X) \not\equiv \exists r.\top$. Without loss of generality, we assume that $C := \sigma(X)$ and $D := \sigma(Y)$ are reduced. Since σ is a unifier of Γ, we have $\exists r.D \sqsubseteq C$. Consequently, Lemma 3.3 yields that C is of the form $C = \exists r.C_1 \sqcap \ldots \sqcap \exists r.C_n$ where $n \geq 1$, $C_1, \ldots, C_n$ are reduced and pairwise incomparable w.r.t. subsumption, and $D \sqsubseteq C_1, \ldots, D \sqsubseteq C_n$.

We use σ to construct a new unifier $\widehat{\sigma}$ as follows:

$$\begin{aligned}
\widehat{\sigma}(X) &:= \exists r.C_1 \sqcap \ldots \sqcap \exists r.C_n \sqcap \exists r.Z \\
\widehat{\sigma}(Y) &:= D \sqcap Z
\end{aligned}$$

where Z is a new variable (i.e., one not occurring in C, D). The second part of Lemma 3.3 implies that $\widehat{\sigma}$ is indeed a unifier of Γ.

Next, we show that $\widehat{\sigma} \precsim \sigma$. To this purpose, we consider the substitution λ that maps Z to C_1, and does not change any of the other variables. Then we have $\lambda(\widehat{\sigma}(X)) = \exists r.C_1 \sqcap \ldots \sqcap \exists r.C_n \sqcap \exists r.C_1 \equiv \exists r.C_1 \sqcap \ldots \sqcap \exists r.C_n = \sigma(X)$ and $\lambda(\widehat{\sigma}(Y)) = D \sqcap C_1 \equiv D = \sigma(Y)$. Note that the second equivalence holds since we have $D \sqsubseteq C_1$.

Since M is complete, there exists a unifier $\theta \in M$ such that $\theta \precsim \widehat{\sigma}$. Transitivity of the relation $\precsim$ thus yields $\theta \precsim \sigma$. Since σ and θ both belong to M, we have completed the proof of the theorem once we have shown that $\sigma \neq \theta$. Assume to the contrary that $\sigma = \theta$. Then we have $\sigma \precsim \widehat{\sigma}$, and thus there exists a substitution μ such that $\mu(\sigma(X)) \equiv \widehat{\sigma}(X)$, i.e.,

$$\exists r.\mu(C_1) \sqcap \ldots \sqcap \exists r.\mu(C_n) \equiv \exists r.C_1 \sqcap \ldots \sqcap \exists r.C_n \sqcap \exists r.Z. \tag{4.1}$$

Recall that the concept terms $C_1, \ldots, C_n$ are reduced and pairwise incomparable w.r.t. subsumption. In addition, since $\sigma(X) = \exists r.C_1 \sqcap \ldots \sqcap \exists r.C_n$ is reduced and not equivalent to $\exists r.\top$, none of the concept terms $C_1, \ldots, C_n$ can be equivalent to $\top$. Finally, Z is a concept name that does not occur in $C_1, \ldots, C_n$. All this implies that $\exists r.C_1 \sqcap \ldots \sqcap \exists r.C_n \sqcap \exists r.Z$ is reduced. Obviously, any reduced form for $\exists r.\mu(C_1) \sqcap \ldots \sqcap \exists r.\mu(C_n)$ is a conjunction of at most n existential restrictions. Thus, Theorem 3.1 shows that the above equivalence (4.1) actually cannot hold.

To sum up, we have shown that M contains two distinct unifiers σ, θ such that $\theta \precsim \sigma$. Since M was an arbitrary complete set of unifiers for Γ, this shows that this unification problem cannot have a minimal complete set of unifiers. $\qquad\square$

5. The decision problem

Before we can describe our decision procedure for $\mathcal{EL}$-unification, we must introduce some notation. An $\mathcal{EL}$-concept term is called an *atom* iff it is a concept name (i.e., concept constant or concept variable) or an existential restriction $\exists r.D$.[7] Obviously, any $\mathcal{EL}$-concept term is (equivalent to) a conjunction of atoms, where the empty conjunction is $\top$. The set $At(C)$ of *atoms of an $\mathcal{EL}$-concept term C* is defined inductively: if $C = \top$, then $At(C) := \emptyset$; if C is a concept name, then $At(C) := \{C\}$; if $C = \exists r.D$ then $At(C) := \{C\} \cup At(D)$; if $C = C_1 \sqcap C_2$, then $At(C) := At(C_1) \cup At(C_2)$.

Concept names and existential restrictions $\exists r.D$ where D is a concept name or $\top$ are called *flat atoms*. An $\mathcal{EL}$-concept term is *flat* iff it is a conjunction of flat atoms (where the

---

[7]Note that $\top$ is *not* an atom.

empty conjunction is $\top$). The $\mathcal{EL}$-unification problem Γ is *flat* iff it consists of equations between flat $\mathcal{EL}$-concept terms. By introducing new concept variables and eliminating $\top$, any $\mathcal{EL}$-unification problem Γ can be transformed in polynomial time into a flat $\mathcal{EL}$-unification problem Γ' such that Γ is solvable iff Γ' is solvable. Thus, we may assume without loss of generality that our input $\mathcal{EL}$-unification problems are flat. Given a flat $\mathcal{EL}$-unification problem $\Gamma = \{C_1 \equiv^? D_1, \dots, C_n \equiv^? D_n\}$, we call the atoms of $C_1, D_1, \dots, C_n, D_n$ the *atoms of* Γ. Atoms of Γ that are not variables (i.e., not elements of N_v) are called *non-variable atoms of* Γ.

The unifier σ of Γ is called *reduced* iff, for all concept variables X occurring in Γ, the $\mathcal{EL}$-concept term $\sigma(X)$ is reduced. It is *ground* iff, for all concept variables X occurring in Γ, the $\mathcal{EL}$-concept term $\sigma(X)$ does not contain variables. Obviously, Γ is solvable iff it has a reduced ground unifier. Given a ground unifier σ of Γ, the *atoms of* σ are the atoms of all the concept terms $\sigma(X)$, where X ranges over all variables occurring in Γ.

Remark 5.1. In the following, we consider situations where all occurrences of a given reduced atom D in a reduced concept term C are replaced by a more general concept term, i.e., by a concept term D' with $D \sqsubseteq D'$. However, when we say *occurrence of D in C*, we mean occurrence modulo equivalence ($\equiv$) rather than syntactic occurrence. For example, if $C = \exists r.(A \sqcap B) \sqcap \exists r.(B \sqcap A)$, $D = \exists r.(A \sqcap B)$, and $D' = \exists r.A$, then the term obtained by replacing all occurrences of D in C by D' should be $\exists r.A \sqcap \exists r.A$, and not $\exists r.A \sqcap \exists r.(B \sqcap A)$. Since C and D are reduced, equivalence is actually the same as being identical up to associativity and commutativity of $\sqcap$. In particular, this means that any concept term that (syntactically) occurs in C and is equivalent to the atom D is also an atom, i.e., only atoms can be replaced by D'. In order to make this meaning of occurrence explicit we will call it *occurrence modulo AC* in the following. We will write $D_1 =_{AC} D_2$ to express that the atoms D_1 and D_2 are identical up to associativity and commutativity of $\sqcap$. Obviously, $D_1 =_{AC} D_2$ implies $D_1 \equiv D_2$.

Lemma 5.2. *Let C, D, D' be $\mathcal{EL}$-concept terms such that D is a reduced atom, $D \sqsubseteq D'$, and C is reduced and contains at least one occurrence of D modulo AC. If C' is obtained from C by replacing all occurrences of D by D', then $C \sqsubseteq C'$.*

Proof. We prove the lemma by induction on the size of C. If $C =_{AC} D$, then $C' = D'$, and thus $C \equiv D \sqsubseteq D' = C'$, which yields $C \sqsubseteq C'$. Thus, assume that $C \neq_{AC} D$. In this case, C cannot be a concept name since it contains the atom D. If $C = \exists r.C_1$, then D occurs in C_1 modulo AC. By induction, we can assume that $C_1 \sqsubseteq C_1'$, where C_1' is obtained from C_1 by replacing all occurrences of D (modulo AC) by D'. Thus, we have $C = \exists r.C_1 \sqsubseteq \exists r.C_1' = C'$ by Corollary 3.2. Finally, assume that $C = C_1 \sqcap \dots \sqcap C_n$ for $n > 1$ atoms $C_1, \dots, C_n$. Since C is reduced, these atoms are incomparable w.r.t. subsumption, and since the atom D occurs in C modulo AC we can assume without loss of generality that D occurs in C_1 modulo AC. Let $C_1', \dots, C_n'$ be respectively obtained from $C_1, \dots, C_n$ by replacing every occurrence of D (modulo AC) by D', and then reducing the concept term obtained this way. By induction, we have $C_1 \sqsubseteq C_1'$. Assume that $C \not\sqsubseteq C'$. Since the concept constructors of $\mathcal{EL}$ are monotone w.r.t. subsumption $\sqsubseteq$, we have $C \sqsubseteq C'$, and thus $C \not\sqsubseteq C'$ means that $C \equiv C'$. Consequently, $C = C_1 \sqcap \dots \sqcap C_n$ and the reduced form of $C_1' \sqcap \dots \sqcap C_n'$ must be equal up to associativity and commutativity of $\sqcap$. If $C_1' \sqcap \dots \sqcap C_n'$ is not reduced, then its reduced form is actually a conjunction of $m < n$ atoms, which contradicts $C \equiv C'$. If $C_1' \sqcap \dots \sqcap C_n'$ is reduced, then $C_1 \sqsubseteq C_1'$ implies that there is an $i \neq 1$ such that

$C_i \equiv C_1'$. However, then $C_i \equiv C_1' \sqsupset C_1$ contradicts the fact that the atoms $C_1, \ldots, C_n$ are incomparable w.r.t. subsumption. $\qquad \square$

Proposition 3.5 says that the inverse strict subsumption order on concept terms is well-founded. We use this fact to obtain a well-founded strict order $\succ$ on ground unifiers.

Definition 5.3. Let σ, θ be ground unifiers of Γ. We define
(1) $\sigma \succeq \theta$ iff $\sigma(X) \sqsubseteq \theta(X)$ holds for all variables X occurring in Γ.
(2) $\sigma \succ \theta$ iff $\sigma \succeq \theta$ and $\theta \not\succeq \sigma$, i.e., iff $\sigma(X) \sqsubseteq \theta(X)$ holds for all variables X occurring in Γ, and $\sigma(X) \sqsubset \theta(X)$ holds for at least one variable X occurring in Γ.

If Γ contains n variables, then $\succeq$ is the n-fold product of the order $\sqsubseteq$ with itself. Since the strict part $\sqsubset$ of the inverse subsumption order $\sqsubseteq$ is well-founded by Proposition 3.5, the strict part $\succ$ of $\succeq$ is also well-founded [13]. The ground unifier σ of Γ is called *is-minimal* iff there is no ground unifier θ of Γ such that $\sigma \succ \theta$. The following proposition is an easy consequence of the fact that $\succ$ is well-founded.

Proposition 5.4. *Let Γ be an $\mathcal{EL}$-unification problem. Then Γ is solvable iff it has an is-minimal reduced ground unifier.*

In the following, we show that is-minimal reduced ground unifiers of flat $\mathcal{EL}$-unification problems satisfy properties that make it easy to check (with an NP-algorithm) whether such a unifier exists or not.

Lemma 5.5. *Let Γ be a flat $\mathcal{EL}$-unification problem and γ an is-minimal reduced ground unifier of Γ. If C is an atom of γ, then there is a non-variable atom D of Γ such that $C \equiv \gamma(D)$.*

The main idea underlying the proof of this crucial lemma is that an atom C of a unifier σ that violates the condition of the lemma (i.e., that is not of the form $C \equiv \gamma(D)$ for a non-variable atom D of Γ) can be replaced by a concept term $\widehat{D}$ such that $C \sqsubset \widehat{D}$, which yields a unifier of Γ that is smaller than σ w.r.t. $\succ$.

Before proving the lemma formally, let us illustrate this idea by two examples.

Example 5.6. First, consider the unification problem
$$\Gamma_1 := \{\exists r.X \sqcap \exists r.A \equiv^? \exists r.X\}.$$
The substitution $\sigma_1 := \{X \mapsto A \sqcap B\}$ is a unifier of Γ_1 that does not satisfy the condition of Lemma 5.5. In fact, B is an atom of σ_1, but none of the non-variable atoms D of Γ_1 (which are A, $\exists r.A$, and $\exists r.X$) satisfy $B \equiv \sigma_1(D)$. The unifier σ_1 is not is-minimal since $\gamma_1 := \{X \mapsto A\}$, which can be obtained from σ_1 by replacing the offending atom B with $\top$, is a unifier of Γ_1 that is smaller than σ_1 w.r.t. $\succ$. The unifier γ_1 is is-minimal, and it clearly satisfies the condition of Lemma 5.5.

Second, consider the unification problem
$$\Gamma_2 := \{X \sqcap \exists r.A \sqcap \exists r.B \equiv^? X\}.$$
The substitution $\sigma_2 := \{X \mapsto \exists r.(A \sqcap B)\}$ is a unifier of Γ_2 that does not satisfy the condition of Lemma 5.5. In fact, $\exists r.(A \sqcap B)$ is an atom of σ_2, but none of the non-variable atoms D of Γ_2 (which are A, B, $\exists r.A$, and $\exists r.B$) satisfy $\exists r.(A \sqcap B) \equiv \sigma_2(D)$. The unifier σ_2 is not is-minimal since $\gamma_2 := \{X \mapsto \exists r.A \sqcap \exists r.B\}$, which can be obtained from σ_2 by replacing the offending atom $\exists r.(A \sqcap B)$ with $\exists r.A \sqcap \exists r.B$, is a unifier of Γ_2 that is

smaller than σ_2 w.r.t. $\succ$. The unifier γ_2 is is-minimal, and it clearly satisfies the condition of Lemma 5.5.

Proof of Lemma 5.5. Assume that γ is an is-minimal reduced ground unifier of Γ. Since γ is reduced, all atoms of γ are reduced. In particular, this implies that C is reduced, and since γ is ground, we know that C is either a concept constant or an existential restriction.

First, assume that C *is of the form A for a concept constant A*, but there is no non-variable atom D of Γ such that $A \equiv \gamma(D)$. This simply means that A does not appear in Γ. Let γ' be the substitution obtained from γ by replacing every occurrence of A by $\top$. Since equivalence in $\mathcal{EL}$ is preserved under replacing concept names by $\top$, and since A does not appear in Γ, it is easy to see that γ' is also a unifier of Γ. However, since $\gamma \succ \gamma'$, this contradicts our assumption that γ is is-minimal.

Second, assume that C *is an existential restriction of the form $\exists r.C_1$*, but there is no non-variable atom D of Γ such that $C \equiv \gamma(D)$. We assume that C is maximal (w.r.t. subsumption) with this property, i.e., for every atom C' of γ with $C \sqsubset C'$, there is a non-variable atom D' of Γ such that $C' \equiv \gamma(D')$. Let $D_1, \ldots, D_\ell$ be all the non-variable atoms of Γ with $C \sqsubseteq \gamma(D_i)$ $(i = 1, \ldots, \ell)$. By our assumptions on C, we actually have $C \sqsubset \gamma(D_i)$ and, by Lemma 3.3, the atom D_i is also an existential restriction $D_i = \exists r.D_i'$ $(i = 1, \ldots, \ell)$. We consider the conjunction

$$\widehat{D} := \gamma(D_1) \sqcap \ldots \sqcap \gamma(D_\ell),$$

which is $\top$ in case $\ell = 0$.

Definition 5.7. Given an $\mathcal{EL}$-concept term F, the concept term $F^{[C/\widehat{D}]}$ is obtained from F by replacing every occurrence of C (modulo AC) by $\widehat{D}$. The substitution $\gamma^{[C/\widehat{D}]}$ is obtained from γ by replacing every occurrence of C (modulo AC) by $\widehat{D}$, i.e., $\gamma^{[C/\widehat{D}]}(X) := \gamma(X)^{[C/\widehat{D}]}$ for all variables X.

We will show in the following that $\gamma^{[C/\widehat{D}]}$ is a unifier of Γ that is smaller than γ w.r.t. $\succ$. This will then again contradict our assumption that γ is is-minimal.

Lemma 5.8. $\gamma \succ \gamma^{[C/\widehat{D}]}$.

Proof. Obviously, $\widehat{D}$ subsumes C. We claim that this subsumption relationship is actually strict. In fact, if $\ell = 0$, then $\widehat{D} = \top$, and since C is an atom, it is not equivalent to $\top$. If $\ell \geq 1$, then $C = \exists r.C_1 \sqsupseteq \exists r.\gamma(D_1') \sqcap \ldots \sqcap \exists r.\gamma(D_\ell')$ would imply (by Corollary 3.2) that there is an $i, 1 \leq i \leq \ell$, with $C_1 \sqsupseteq \gamma(D_i')$. However, this would yield $C = \exists r.C_1 \sqsupseteq \exists r.\gamma(D_i') = \gamma(D_i)$, which contradicts the fact that $C \sqsubset \gamma(D_i)$. Thus, we have shown that $C \sqsubset \widehat{D}$. Lemma 5.2 implies that $\gamma \succ \gamma'$. $\square$

To complete the proof of Lemma 5.5, it remains to show the next lemma.

Lemma 5.9. $\gamma^{[C/\widehat{D}]}$ *is a unifier of* Γ.

Proof. Consider an equation in Γ of the form $L_1 \sqcap \ldots \sqcap L_m \equiv^? R_1 \sqcap \ldots \sqcap R_n$ where $L_1, \ldots, L_m$ and $R_1, \ldots, R_n$ are flat atoms, and define $L := \gamma(L_1 \sqcap \ldots \sqcap L_m)$ and $R := \gamma(R_1 \sqcap \ldots \sqcap R_n)$. We know that L, R are conjunctions of atoms of the form $L = A_1 \sqcap \ldots \sqcap A_\mu$ and $R = B_1 \sqcap \ldots \sqcap B_\nu$, where each conjunct $A_1, \ldots, A_\mu, B_1, \ldots, B_\nu$ is a reduced ground atom that is either an atom of γ or equal to $\gamma(E)$ for a non-variable atom E of Γ. Since γ is a unifier of Γ, we have $L \equiv R$.

(1) Since C is an atom, we obviously have $L^{[C/\widehat{D}]} = A_1^{[C/\widehat{D}]} \sqcap \ldots \sqcap A_\mu^{[C/\widehat{D}]}$ and $R^{[C/\widehat{D}]} = B_1^{[C/\widehat{D}]} \sqcap \ldots \sqcap B_\nu^{[C/\widehat{D}]}$. Now, we show that $L^{[C/\widehat{D}]} = \gamma^{[C/\widehat{D}]}(L_1 \sqcap \ldots \sqcap L_m)$ and $R^{[C/\widehat{D}]} = \gamma^{[C/\widehat{D}]}(R_1 \sqcap \ldots \sqcap R_n)$. We concentrate on proving the first identity since the second one can be shown analogously. To show the first identity, it is enough to prove that $\gamma(L_j)^{[C/\widehat{D}]} = \gamma^{[C/\widehat{D}]}(L_j)$ holds for all $j, 1 \le j \le m$.

(a) If L_j is a variable X, then $\gamma^{[C/\widehat{D}]}(X) = \gamma(X)^{[C/\widehat{D}]}$ holds by the definition of $\gamma^{[C/\widehat{D}]}$.

(b) If L_j is a concept constant A, then $A^{[C/\widehat{D}]} = A$ since C is an existential restriction. Thus, we have $\gamma^{[C/\widehat{D}]}(A) = A = A^{[C/\widehat{D}]} = \gamma(A)^{[C/\widehat{D}]}$.

(c) Otherwise, L_j is an existential restriction $\exists r_j.L_j'$. By our assumption on C, we have $C \not\equiv \gamma(L_j)$, and thus $\gamma(L_j)^{[C/\widehat{D}]} = \exists r_j.\left(\gamma(L_j)^{[C/\widehat{D}]}\right)$. In addition, we have $\gamma^{[C/\widehat{D}]}(L_j) = \exists r_j.\gamma^{[C/\widehat{D}]}(L_j')$. Thus, it is enough to show $\gamma(L_j')^{[C/\widehat{D}]} = \gamma^{[C/\widehat{D}]}(L_j')$. Since L_j is a flat atom, we know that L_j' is either a concept constant, the top-concept $\top$, or a concept variable. In the first to cases, we can show $\gamma(L_j')^{[C/\widehat{D}]} = \gamma^{[C/\widehat{D}]}(L_j')$ as in (1b), and in the third case we can show this identity as in (1a).

(2) Because of (1), if we can prove that $L^{[C/\widehat{D}]} \equiv R^{[C/\widehat{D}]}$, then we have shown that $\gamma^{[C/\widehat{D}]}$ solves the equation $L_1 \sqcap \ldots \sqcap L_m \equiv^? R_1 \sqcap \ldots \sqcap R_n$.

Without loss of generality, we concentrate on showing that $L^{[C/\widehat{D}]} \sqsubseteq R^{[C/\widehat{D}]}$. Since $L^{[C/\widehat{D}]} = A_1^{[C/\widehat{D}]} \sqcap \ldots \sqcap A_\mu^{[C/\widehat{D}]}$ and $R^{[C/\widehat{D}]} = B_1^{[C/\widehat{D}]} \sqcap \ldots \sqcap B_\nu^{[C/\widehat{D}]}$, it is thus sufficient to show that, for every $i, 1 \le i \le \nu$, there exists a $j, 1 \le j \le \mu$, such that $A_j^{[C/\widehat{D}]} \sqsubseteq B_i^{[C/\widehat{D}]}$ (see (3) of Lemma 3.4). Since $L = A_1 \sqcap \ldots \sqcap A_\mu \sqsubseteq B_1 \sqcap \ldots \sqcap B_\nu = R$ and $A_1, \ldots, A_\mu, B_1, \ldots, B_\nu$ are atoms, we actually know that, for every $i, 1 \le i \le \nu$, there exists a $j, 1 \le j \le \mu$, such that $A_j \sqsubseteq B_i$. Thus, it is sufficient to show that $A_j \sqsubseteq B_i$ implies $A_j^{[C/\widehat{D}]} \sqsubseteq B_i^{[C/\widehat{D}]}$. This is an easy consequence of the next lemma since A_i, B_j satisfy the conditions of this lemma. $\square$

Lemma 5.10. *Let A, B be reduced ground atoms such that B is an atom of γ or of the form $\gamma(D)$ for a non-variable atom D of Γ. If $A \sqsubseteq B$, then $A^{[C/\widehat{D}]} \sqsubseteq B^{[C/\widehat{D}]}$.*

Proof. We show $A^{[C/\widehat{D}]} \sqsubseteq B^{[C/\widehat{D}]}$ by induction on the size of A.

(1) First, assume that $A =_{AC} C$, which implies that $A^{[C/\widehat{D}]} = \widehat{D} = \gamma(D_1) \sqcap \ldots \sqcap \gamma(D_n)$.

(a) If B is of the form $B \equiv \gamma(D)$ for a non-variable atom D of Γ, then there is an $h, 1 \le h \le n$, such that $D = D_h$, which shows that $A^{[C/\widehat{D}]} \sqsubseteq B$. Since $C \sqsubseteq \widehat{D}$ and the constructors of $\mathcal{EL}$ are monotone w.r.t. subsumption, we also have $B \sqsubseteq B^{[C/\widehat{D}]}$, and thus $A^{[C/\widehat{D}]} \sqsubseteq B^{[C/\widehat{D}]}$.

(b) Assume that B is an atom of γ. If $B =_{AC} C$, then $B^{[C/\widehat{D}]} = \widehat{D}$, and thus $A^{[C/\widehat{D}]} = B^{[C/\widehat{D}]}$, which implies $A^{[C/\widehat{D}]} \sqsubseteq B^{[C/\widehat{D}]}$. Otherwise, since C, B are reduced atoms, $B \neq_{AC} C$ implies $B \not\equiv C$. Together with $C \equiv A \sqsubseteq B$, this shows that $C \sqsubset B$. Thus, the maximality of C implies that there is a non-variable atom D of Γ such that $B \equiv \gamma(D)$. Thus, we are actually in case (a), which yields $A^{[C/\widehat{D}]} \sqsubseteq B^{[C/\widehat{D}]}$.

(2) Now, assume that $A \neq_{AC} C$. If there is no occurrence (modulo AC) of C in A, then we have $A^{[C/\widehat{D}]} = A \sqsubseteq B \sqsubseteq B^{[C/\widehat{D}]}$.

Otherwise, A is of the form $A = \exists s.E$ and C occurs in E (modulo AC). Obviously, $A \sqsubseteq B$ then implies that B is of the form $B = \exists s.F$ with $E \sqsubseteq F$. The concept terms E, F are conjunctions of reduced ground atoms, i.e., $E = E_1 \sqcap \ldots \sqcap E_\kappa$ and $F = F_1 \sqcap \ldots \sqcap F_\lambda$ where $E_1, \ldots, E_\kappa, F_1, \ldots, F_\lambda$ are reduced ground atoms. By Corollary 3.2, for every $h, 1 \le h \le \lambda$, there exists $k, 1 \le k \le \kappa$ such that $E_k \sqsubseteq F_h$.

In order to be able to assume, by induction, that $E_k \sqsubseteq F_h$ implies $E_k^{[C/\widehat{D}]} \sqsubseteq F_h^{[C/\widehat{D}]}$, we must show that the conditions in the statement of the lemma hold for the concept terms E_k, F_h, where E_k plays the rôle of A and F_h plays the rôle of B. Since we already know that $E_1, \ldots, E_\kappa, F_1, \ldots, F_\lambda$ are reduced ground atoms, it is sufficient to show that each of the atoms $F_1, \ldots, F_\lambda$ is an atom of γ or of the form $\gamma(D)$ for a non-variable atom D of Γ. We know that $B = \exists s.(F_1 \sqcap \ldots \sqcap F_\lambda)$ is an atom of γ or an instance (w.r.t. γ) of a non-variable atom of Γ. In the first case, the atoms $F_1, \ldots, F_\lambda$ are clearly also atoms of γ. In the second case, $B = \gamma(D')$ for a non-variable atom D' of Γ. If D' is a ground atom, then $F_1, \ldots, F_\lambda$ are also ground atoms that are atoms of Γ, and thus they are instances (w.r.t. γ) of non-variable atoms of Γ. Otherwise, since Γ is flat, D' is of the form $\exists s.X$ for a variable X and $\gamma(X) = F_1 \sqcap \ldots \sqcap F_\lambda$. In this case, $F_1, \ldots, F_\lambda$ are clearly atoms of γ.

Thus, we can assume by induction:

($*$) for every $h, 1 \le h \le \lambda$, there exists $k, 1 \le k \le \kappa$ such that $E_k^{[C/\widehat{D}]} \sqsubseteq F_h^{[C/\widehat{D}]}$

It remains to show that this implies $A^{[C/\widehat{D}]} \sqsubseteq B^{[C/\widehat{D}]}$.

(a) If $B \ne_{AC} C$, then $A^{[C/\widehat{D}]} = \exists s.(E_1^{[C/\widehat{D}]} \sqcap \ldots \sqcap E_\kappa^{[C/\widehat{D}]})$ and $B^{[C/\widehat{D}]} = \exists s.(F_1^{[C/\widehat{D}]} \sqcap \ldots \sqcap F_\lambda^{[C/\widehat{D}]})$, and thus property ($*$) yields $A^{[C/\widehat{D}]} \sqsubseteq B^{[C/\widehat{D}]}$.

(b) Assume that $B =_{AC} C$. In this case, C cannot occur (modulo AC) in any of the concept terms $F_1, \ldots, F_h$, which implies that $B = \exists s.(F_1 \sqcap \ldots \sqcap F_\lambda) = \exists s.(F_1^{[C/\widehat{D}]} \sqcap \ldots \sqcap F_\lambda^{[C/\widehat{D}]})$. Since we have $A^{[C/\widehat{D}]} = \exists s.(E_1^{[C/\widehat{D}]} \sqcap \ldots \sqcap E_\kappa^{[C/\widehat{D}]})$, property ($*$) yields $A^{[C/\widehat{D}]} \sqsubseteq B$. Since we also have $B \sqsubseteq B^{[C/\widehat{D}]}$, this yields $A^{[C/\widehat{D}]} \sqsubseteq B^{[C/\widehat{D}]}$.

Thus, we have shown in all cases that $A^{[C/\widehat{D}]} \sqsubseteq B^{[C/\widehat{D}]}$, which completes the proof of Lemma 5.10. □

Overall, we have thus completed the proof of Lemma 5.5. The next proposition is an easy consequence of this lemma.

Proposition 5.11. *Let Γ be a flat $\mathcal{EL}$-unification problem and γ an is-minimal reduced ground unifier of Γ. If X is a concept variable occurring in Γ, then $\gamma(X) \equiv \top$ or there are non-variable atoms $D_1, \ldots, D_n$ ($n \ge 1$) of Γ such that $\gamma(X) \equiv \gamma(D_1) \sqcap \ldots \sqcap \gamma(D_n)$.*

Proof. If $\gamma(X) \not\equiv \top$, then it is a non-empty conjunction of atoms, i.e., there are atoms $C_1, \ldots, C_n$ ($n \ge 1$) such that $\gamma(X) = C_1 \sqcap \ldots \sqcap C_n$. Then $C_1, \ldots, C_n$ are atoms of γ, and thus Lemma 5.5 yields non-variable atoms $D_1, \ldots, D_n$ of Γ such that $C_i \equiv \gamma(D_i)$ for $i = 1, \ldots n$. Consequently, $\gamma(X) \equiv \gamma(D_1) \sqcap \ldots \sqcap \gamma(D_n)$. □

This proposition suggests the following non-deterministic algorithm for deciding solvability of a given flat $\mathcal{EL}$-unification problem.

Algorithm 5.12. Let Γ be a flat $\mathcal{EL}$-unification problem.

(1) For every variable X occurring in Γ, guess a finite, possibly empty, set S_X of non-variable atoms of Γ.
(2) We say that the variable X *directly depends on* the variable Y if Y occurs in an atom of S_X. Let *depends on* be the transitive closure of *directly depends on*. If there is a variable that depends on itself, then the algorithm returns "fail." Otherwise, there exists a strict linear order $>$ on the variables occurring in Γ such that $X > Y$ if X depends on Y.
(3) We define the substitution σ along the linear order $>$:
 - If X is the least variable w.r.t. $>$, then S_X does not contain any variables. We define $\sigma(X)$ to be the conjunction of the elements of S_X, where the empty conjunction is $\top$.
 - Assume that $\sigma(Y)$ is defined for all variables $Y < X$. Then S_X only contains variables Y for which $\sigma(Y)$ is already defined. If S_X is empty, then we define $\sigma(X) := \top$. Otherwise, let $S_X = \{D_1, \ldots, D_n\}$. We define $\sigma(X) := \sigma(D_1) \sqcap \ldots \sqcap \sigma(D_n)$.
(4) Test whether the substitution σ computed in the previous step is a unifier of Γ. If this is the case, then return σ; otherwise, return "fail."

This algorithm is trivially *sound* since it only returns substitutions that are unifiers of Γ. In addition, it obviously always terminates. Thus, to show correctness of our algorithm, it is sufficient to show that it is complete.

Lemma 5.13 (Completeness). *If Γ is solvable, then there is a way of guessing in Step 1 subsets S_X of the non-variable atoms of Γ such that the* depends on *relation determined in Step 2 is acyclic and the substitution σ computed in Step 3 is a unifier of Γ.*

Proof. If Γ is solvable, then it has an is-minimal reduced ground unifier γ. By Proposition 5.11, for every variable X occurring in Γ we have $\gamma(X) \equiv \top$ or there are non-variable atoms $D_1, \ldots, D_n$ ($n \geq 1$) of Γ such that $\gamma(X) \equiv \gamma(D_1) \sqcap \ldots \sqcap \gamma(D_n)$. If $\gamma(X) \equiv \top$, then we define $S_X := \emptyset$. Otherwise, we define $S_X := \{D_1, \ldots, D_n\}$.

We show that the relation *depends on* induced by these sets S_X is acyclic, i.e., there is no variable X such that X depends on itself. If X directly depends on Y, then Y occurs in an element of S_X. Since S_X consists of non-variable atoms of the flat unification problem Γ, this means that there is a role name r such that $\exists r.Y \in S_X$. Consequently, we have $\gamma(X) \sqsubseteq \exists r.\gamma(Y)$. Thus, if X depends on X, then there are $k \geq 1$ role names $r_1, \ldots, r_k$ such that $\gamma(X) \sqsubseteq \exists r_1. \cdots \exists r_k.\gamma(X)$. This is clearly not possible since $\gamma(X)$ cannot be subsumed by an $\mathcal{EL}$-concept term whose role depth is larger than the role depth of $\gamma(X)$.

To show that the substitution σ induced by the sets S_X is a unifier of Γ, we prove that σ is equivalent to γ, i.e., $\sigma(X) \equiv \gamma(X)$ holds for all variables X occurring in Γ. The substitution σ is defined along the linear order $>$. If X is the least variable w.r.t. $>$, then the elements of S_X do not contain any variables. If S_X is empty, then $\sigma(X) = \top \equiv \gamma(X)$. Otherwise, let $S_X = \{D_1, \ldots, D_n\}$. Since the atoms D_i do not contain variables, we have $D_i = \gamma(D_i)$. Thus, the definitions of S_X and of σ yield $\sigma(X) = D_1 \sqcap \ldots \sqcap D_n = \gamma(D_1) \sqcap \ldots \sqcap \gamma(D_n) \equiv \gamma(X)$.

Assume that $\sigma(Y) \equiv \gamma(Y)$ holds for all variables $Y < X$. If $S_X = \emptyset$, then we have again $\sigma(X) = \top \equiv \gamma(X)$. Otherwise, let $S_X = \{D_1, \ldots, D_n\}$. Since the atoms D_i contain only variables that are smaller than X, we have $\sigma(D_i) \equiv \gamma(D_i)$ by induction. Thus, the definitions of S_X and of σ yield $\sigma(X) = \sigma(D_1) \sqcap \ldots \sqcap \sigma(D_n) \equiv \gamma(D_1) \sqcap \ldots \sqcap \gamma(D_n) \equiv \gamma(X)$. $\square$

Note that our proof of completeness actually shows that, up to equivalence, the algorithm returns all is-minimal reduced ground unifiers of Γ.

Theorem 5.14. $\mathcal{EL}$-*unification is NP-complete.*

Proof. NP-hardness follows from the fact that $\mathcal{EL}$-matching is NP-complete [24].[8] To show that the problem can be decided by a non-deterministic polynomial-time algorithm, we analyze the complexity of our algorithm. Obviously, guessing the sets S_X (Step 1) can be done within NP. Computing the *depends on* relation and checking it for acyclicity (Step 2) is clearly polynomial.

Steps 3 and 4 are more problematic. In fact, since a variable may occur in different atoms of Γ, the substitution σ computed in Step 3 may be of exponential size. This is actually the same reason that makes a naive algorithm for syntactic unification compute an exponentially large most general unifier [16]. As in the case of syntactic unification, the solution to this problem is basically structure sharing. Instead of computing the substitution σ explicitly, we view its definition as an acyclic TBox. To be more precise, for every concept variable X occurring in Γ, the TBox $\mathcal{T}_\sigma$ contains the concept definition $X \doteq \top$ if $S_X = \emptyset$ and $X \doteq D_1 \sqcap \ldots \sqcap D_n$ if $S_X = \{D_1, \ldots, D_n\}$ ($n \geq 1$). Instead of computing σ in Step 3, we compute $\mathcal{T}_\sigma$. Because of the acyclicity test in Step 2, we know that $\mathcal{T}_\sigma$ is an acyclic TBox. The size of $\mathcal{T}_\sigma$ is obviously polynomial in the size of Γ, and thus this modified Step 3 is polynomial.

It is easy to see that applying the substitution σ to a concept term C is the same as expanding C w.r.t. the TBox $\mathcal{T}_\sigma$, i.e., $\sigma(C) = C^{\mathcal{T}_\sigma}$. This implies that, for every equation $C \equiv^? D$ in Γ, we have $C \equiv_{\mathcal{T}_\sigma} D$ iff $\sigma(C) \equiv \sigma(D)$. Thus, testing in Step 4 whether σ is a unifier of Γ can be reduced to testing whether $C \equiv_{\mathcal{T}_\sigma} D$ holds for every equation $C \equiv^? D$ in Γ. Since subsumption (and thus equivalence) in $\mathcal{EL}$ w.r.t. acyclic TBoxes can be decided in polynomial time [4], this completes the proof of the theorem. $\square$

In Subsection 2.3, we have shown that there exists a polynomial-time reduction of unification modulo an acyclic TBox to unification without a TBox. Thus, Theorem 5.14 also yields the exact complexity for $\mathcal{EL}$-unification w.r.t. acyclic TBoxes.

Corollary 5.15. *$\mathcal{EL}$-unification w.r.t. acyclic TBoxes is NP-complete.*

Proof. The problem is in NP since Theorem 2.8 states that there is a polynomial-time reduction of $\mathcal{EL}$-unification w.r.t. acyclic TBoxes to $\mathcal{EL}$-unification, and we have just shown that $\mathcal{EL}$-unification is in NP.

NP-hardness for $\mathcal{EL}$-unification w.r.t. acyclic TBoxes follows from NP-hardness of $\mathcal{EL}$-unification since $\mathcal{EL}$-unification can be viewed as the special case of $\mathcal{EL}$-unification w.r.t. acyclic TBoxes where the TBox is empty. $\square$

6. A GOAL-ORIENTED ALGORITHM

The NP-algorithm introduced in the previous section is a typical "guess and then test" NP-algorithm, and thus it is unlikely that a direct implementation of this algorithm will perform well in practice. Here, we introduce a more goal-oriented unification algorithm for $\mathcal{EL}$, in which non-deterministic decisions are only made if they are triggered by "unsolved parts" of the unification problem.

As in the previous section, we assume without loss of generality that our input unification problem Γ_0 is flat. For a given flat equation $C \equiv^? D$, the concept terms C, D are thus

---

[8]The NP-hardness proof in [24] is by reduction of SAT. This reduction employs two concept constants and four role names. However, the roles are mainly used to encode several (matching) equations into a single one. When using a set of equations rather than a single equation, one role name is sufficient.

conjunctions of flat atoms. We will often view such an equation as consisting of four sets: the left-hand side C is given by the set of variables occurring in the top-level conjunction of C, together with the set of non-variable atoms occurring in this top-level conjunction; the right-hand side D is given by the set of variables occurring in the top-level conjunction of D, together with the set of non-variable atoms occurring in this top-level conjunction. To be more precise, let e denote the equation $C \equiv^? D$, where $C = X_1 \sqcap \ldots \sqcap X_m \sqcap A_1 \sqcap \ldots \sqcap A_k$ and $D = Y_1 \sqcap \ldots \sqcap Y_n \sqcap B_1 \sqcap \ldots \sqcap B_\ell$ for concept variables $X_1, \ldots, X_m, Y_1, \ldots, Y_n$ and non-variable atoms $A_1, \ldots, A_k, B_1, \ldots B_\ell$. Then we define

$$LVar(e) := \{X_1, \ldots, X_m\}, \quad RVar(e) := \{Y_1, \ldots, Y_n\},$$
$$LAto(e) := \{A_1, \ldots, A_k\}, \quad RAto(e) := \{B_1, \ldots B_\ell\}.$$

Obviously, the equation $e : C \equiv^? D$ is uniquely determined (up to associativity, commutativity, and idempotency of conjunction) by the four sets $LVar(e), LAto(e), RVar(e), RAto(e)$. Instead of viewing an equation e as being given by a pair of concept terms, we can thus also view it as being given by these four sets. In the following, it will often be convenient to employ this representation of equations. If, with this point of view, we say that we *add an atom* to the set $LAto(e)$ or $RAto(e)$, then this means, for the other point of view, that we conjoin this atom to the top-level conjunction of the left-hand side or right-hand side of the equation. In addition, if we say that the equation e *contains* the variable X, then we mean that $X \in LVar(e) \cup RVar(e)$. Similarly, if we say that the left-hand side of e contains X, then we mean that $X \in LVar(e)$, and if we say that the right-hand side of e contains X, then we mean that $X \in RVar(e)$).[9]

In addition to the unification problem itself, the algorithm also maintains, for every variable X occurring in the input problem Γ_0, a set S_X of non-variable atoms of Γ_0. Initially, all the sets S_X are empty. We call the set S_X the current assignment for X, and the collection of all these sets the *current assignment*. Throughout the run of our goal-oriented algorithm, we will ensure that the current assignment is *acyclic* in the sense that no variable depends on itself w.r.t. this assignment (see (2) of Algorithm 5.12). An acyclic assignment induces a substitution σ, as defined in (3) of Algorithm 5.12. We call this substitution the *current substitution*. Initially, the current substitution maps all variables to $\top$.

The algorithm applies rules that can

(1) change an equation of the unification problem by adding non-variable atoms of the input problem Γ_0 to one side of the equation;

(2) introduce a new flat equation of the form $C \sqcap B \equiv B$, where C, B are atoms of the input problem Γ_0 or $\top$;

(3) add non-variable atoms of the input problem Γ_0 to the sets S_X.

Another property that is maintained throughout the run of our algorithm is that all equations e are *expanded* w.r.t. the current assignment in the following sense: for all variables X we have

$$A \in S_X \wedge X \in LVar(e) \Rightarrow A \in LAto(e) \quad \text{and} \quad A \in S_X \wedge X \in RVar(e) \Rightarrow A \in RAto(e).$$

Given a flat equation e that contains the variable X, the *expansion* of e w.r.t. the assignment S_X for X is defined as follows: if $X \in LVar(e)$ then all elements of S_X are added to $LAto(e)$, and if $X \in RVar(e)$ then all elements of S_X are added to $RAto(e)$.

---

[9] Note that occurrences of X inside non-variable atoms $\exists r.X \in LAto(e) \cup RAto(e)$ are not taken into consideration here.

The L-variant of the Eager-Assignment rule applies to the equation e if there is an unfinished variable $X \in LVar(e)$ such that

- all variables $Z \in (LVar(e) \setminus \{X\}) \cup RVar(e)$ are finished;
- $LAto(e) = \emptyset$.

Its application sets $S_X := RAto(e)$.

    (1) If this makes the current assignment cyclic, then return "fail."

    (2) Otherwise, label X as finished and expand all equations containing X w.r.t. the new assignment for X.
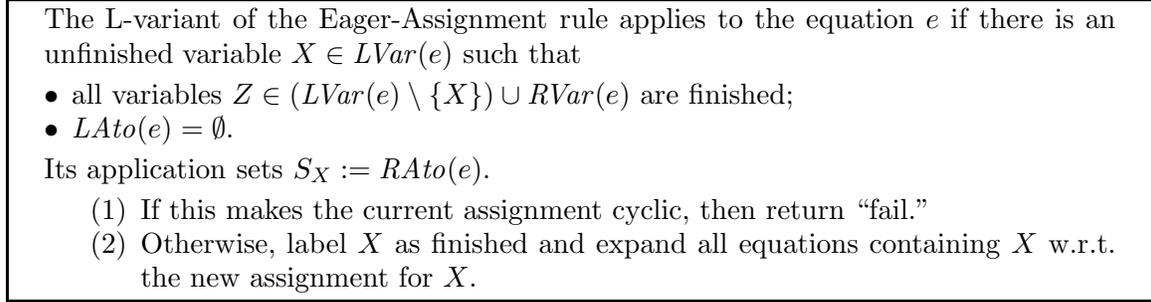

Figure 1: The Eager-Assignment rule in its L-variant. The R-variant is obtained by exchanging the rôles of the two sides of the equation.

The following lemma is an immediate consequence of the definition of expanded equations and of the construction of the current substitution.

Lemma 6.1. *If the equation $C \equiv^? D$ is expanded w.r.t. the current assignment, then $LAto(C \equiv^? D) = RAto(C \equiv^? D)$ implies that the current substitution σ solves this equation, i.e., $\sigma(C) \equiv \sigma(D)$.*

We say that an equation e is *solved* if $LAto(e) = RAto(e)$. An atom $A \in LAto(e) \cap RAto(e)$ is called *solved in e*; atoms $A \in LAto(e) \cup RAto(e)$ that are not solved in e are called *unsolved in e*. Obviously, an equation e is solved iff all atoms $A \in LAto(e) \cup RAto(e)$ are solved in e.

Basically, in each step, the goal-oriented algorithm considers an unsolved equation and an unsolved atom in this equation, and tries to solve it. Picking the unsolved equation and the unsolved atom in it is don't care non-deterministic, i.e., there is no need to backtrack over such a choice. Once an unsolved equation and an unsolved atom in it was picked, don't know non-determinism comes in since there may be several possibilities for how to solve this atom in the equation, some of which may lead to overall success whereas others won't. In some cases, however, a given equation uniquely determines the assignment for a certain variable X. In this case, we make this assignment and then label the variable X as *finished*. This has the effect that the set S_X can no longer be extended. Initially, none of the variables occurring in the input unification problem is labeled as finished. We say that the variable X is *unfinished* if it is not labeled as finished.

Algorithm 6.2. Let Γ_0 be a flat $\mathcal{EL}$-unification problem. We define $\Gamma := \Gamma_0$ and $S_X := \emptyset$ for all variables X occurring in Γ_0. None of these variables is labeled as finished.

As long as Γ contains an unsolved equation, do the following:

(1) If the Eager-Assignment rule applies to some equation e, then apply it to this equation (see Figure 1).

(2) Otherwise, let e be an unsolved equation and A an unsolved atom in e. If neither of the rules Decomposition (see Figure 2) and Extension (see Figure 3) applies to A in e, then return "fail." If one of these rules applies to A in e, then (don't know) non-deterministically choose one of these rules and apply it.

Once all equations of Γ are solved, return the substitution σ that is induced by the current assignment.

The *Eager-Assignment rule* is described in Figure 1. Note that, after a non-failing application of this rule, the equation it was applied to is solved since the expansion of this

---

The L-variant of the Decomposition rule applies to the unsolved atom A in the equation e if

- $A \in LAto(e) \setminus RAto(e)$;
- A is of the form $A = \exists r.C$;
- there is at least one atom of the form $\exists r.B \in RAto(e)$.

Its application chooses (don't know) non-deterministically an atom of the form $\exists r.B \in RAto(e)$ and

- adds $\exists r.C$ to $RAto(e)$;
- creates a new equation $C \sqcap B \equiv^? B$ and expands it w.r.t. the assignments of all variables contained in this equation, unless this equation has already been generated before. If the equation has already been generated before, it is not generated again.

Figure 2: The Decomposition rule in its L-variant. The R-variant is obtained by exchanging the rôles of the two sides of the equation.

equation w.r.t. the new assignment for X adds all elements of $RAto(e)$ to $LAto(e)$. As an example, consider the equations

$$Y \equiv^? \top, \quad Z \equiv^? \exists r.\top, \quad X \sqcap Y \equiv^? Z,$$

and assume that $S_X = S_Y = S_Z = \emptyset$ and none of the three variables X, Y, Z is finished. An application of the Eager-Assignment rule to the first equation labels Y as finished, but does not change anything else. The subsequent application of the Eager-Assignment rule to the second equation changes the assignment for Z to $S_Z = \{\exists r.\top\}$, labels Z as finished, and expands the second and the third equation w.r.t. the new assignment for Z. Thus, we now have the equations

$$Y \equiv^? \top, \quad Z \sqcap \exists r.\top \equiv^? \exists r.\top, \quad X \sqcap Y \equiv^? Z \sqcap \exists r.\top.$$

Since Y, Z are finished, the Eager-Assignment rule can now be applied to the third equation. This changes the assignment for X to $S_X = \{\exists r.\top\}$, labels X as finished, and adds $\exists r.\top$ to the left-hand side of the third equation. Now all equations are solved. The current assignment induces a substitution σ with $\sigma(X) = \exists r.\top = \sigma(Z)$ and $\sigma(Y) = \top$, which is a unifier of the original set of equations.

The *Decomposition rule* is described in Figure 2. This rule solves the unsolved atom $A = \exists r.C$ by adding it to the other side. For this to be admissible, one needs a more specific atom $\exists r.B$ on that side, where the "more specific" is meant to hold after application of the unifier. Thus, to ensure that the unifier σ computed by the algorithm satisfies $\sigma(\exists r.B) \sqsubseteq \sigma(\exists r.C)$, the rule adds the new equation $C \sqcap B \equiv^? B$. Obviously, if the substitution σ solves this equation, then it satisfies $\sigma(B) \sqsubseteq \sigma(C)$, and thus $\sigma(\exists r.B) \sqsubseteq \sigma(\exists r.C)$. As an example, consider the equation

$$\exists r.X \sqcap \exists r.A \equiv^? \exists r.A,$$

and assume that $S_X = \emptyset$ and that X is unfinished. An application of the L-variant of the Decomposition rule to this equation adds $\exists r.X$ to the right-hand side of this equation, and thus solves it. In addition, it generates the new equation $X \sqcap A \equiv^? A$, which is solved. The current assignment induces a substitution σ with $\sigma(X) = \top$, which solves the original equation.
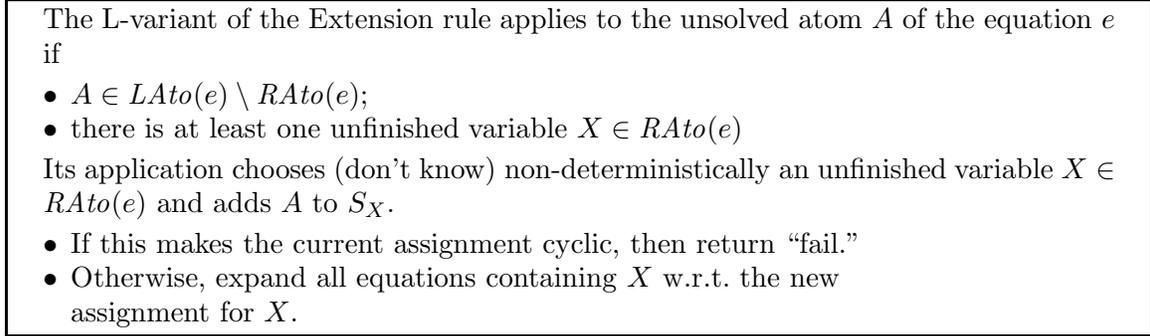
> The L-variant of the Extension rule applies to the unsolved atom A of the equation e if
>
> - $A \in LAto(e) \setminus RAto(e)$;
> - there is at least one unfinished variable $X \in RAto(e)$
>
> Its application chooses (don't know) non-deterministically an unfinished variable $X \in RAto(e)$ and adds A to S_X.
>
> - If this makes the current assignment cyclic, then return "fail."
> - Otherwise, expand all equations containing X w.r.t. the new assignment for X.

Figure 3: The Extension rule in its L-variant. The R-variant is obtained by exchanging the rôles of the two sides of the equation.

The *Extension rule* is described in Figure 3. Basically, this rule solves the unsolved atom A by extending with this atom the assignment of an unfinished variable contained in the other side of the equation. As an example, consider the equation

$$A \sqcap \exists r.\top \equiv^? \exists r.\top \sqcap X,$$

where A is a concept constant, $S_X = \emptyset$, and X is unfinished. An application of the Extension rule to A in this equation extends the assignment for X to $S_X = \{A\}$, and expands this equation by adding A to the right-hand side. The equation obtained this way is solved. The substitution σ induced by the current assignment replaces X by A, and solves the original equation.

Theorem 6.3. *Algorithm 6.2 is an NP-algorithm for testing solvability of flat $\mathcal{EL}$-unification problems.*

First, we show that the algorithm is indeed an NP-algorithm. For this, we consider all *runs* of the algorithm, where for every (don't care) non-deterministic choice exactly one alternative is taken. Since a single rule application can obviously be realized in polynomial time, it is sufficient to show the following lemma.

Lemma 6.4 (Termination). *Every run of the algorithm terminates after a polynomial number of rule applications.*

Proof. Each application of the Eager-Assignment rule finishes an unfinished variable. Thus, since finished variables never become unfinished again, it can only be applied k times, where k is the number of variables occurring in the input unification problem Γ_0. This number is clearly linearly bounded by the size of Γ_0.

Every application of the Decomposition rule or the Extension rule turns an unsolved atom in an equation into a solved one, and a solved atom in an equation never becomes unsolved again in this equation. For a fixed equation, in the worst case every atom of Γ_0 may become an unsolved atom of the equation that needs to be solved. There is, however, only a linear number of atoms of Γ_0. Each equation considered during the run of the algorithm is either descended from an original equation of Γ_0, or from an equation of the form $C \sqcap B \equiv^? B$ for atoms $\exists r.B$ and $\exists r.C$ of Γ_0. Thus, the number of equations is also polynomially bounded by the size of Γ_0. Overall, this shows that the Decomposition rule and the Extension rule can only be applied a polynomial number of times. $\square$

Next, we show soundness of Algorithm 6.2. We call a run of this algorithm *non-failing* if it terminates with a unification problem containing only solved equations.

Lemma 6.5 (Soundness). *Let Γ_0 be a flat $\mathcal{EL}$-unification problem. The substitution σ returned after a successful run of Algorithm 6.2 on input Γ_0 is an $\mathcal{EL}$-unifier of Γ_0.*

Proof. First, note that the rules employed by Algorithm 6.2 indeed preserve the two invariants mentioned before:

(1) the current assignment is always acyclic;
(2) all equations are expanded.

In fact, whenever the current assignment is extended, the rules test acyclicity (and return "fail," if it is not satisfied). In addition, they expand all equations w.r.t. the new assignment.

Now, assume that the run of the algorithm has terminated with the $\mathcal{EL}$-unification problem $\widehat{\Gamma}$, in which all equations are solved. The first invariant ensures that the final assignment constructed by the run is acyclic, and thus indeed induces a substitution σ. Because of the second invariant, Lemma 6.1 applies, and thus we know that σ is a solution of $\widehat{\Gamma}$.

It remains to show that the substitution σ is also a solution of the input problem Γ_0. To this purpose, we take all the equations that were considered during the run of the algorithm, i.e., present in Γ_0 or in any of the other unification problems generated during the run. Let $\mathcal{E}$ denote the set of these equations. We define the relation $\to$ on $\mathcal{E}$ as follows: $e \to e'$ if e was transformed into e' using one of the rules of Algorithm 6.2. To be more precise, the Eager-Assignment rule transforms equations containing X from the current unification problem Γ by expanding them w.r.t. the new assignment for X. The same is true for the Extension rule. The decomposition rule transforms an equation e containing the unsolved atom $A = \exists r.C$ by adding this atom to the other side, which needs to contain an atom of the form $\exists r.B$. For this new equation e', we have $e \to e'$. The decomposition rule may also generate a new equation e'' of the form $C \sqcap B \equiv^? B$ (if this equation was not generated before). However, we do not view this equation as a successor of e w.r.t. $\to$, i.e., we do *not* have $e \to e''$. Equations $C \sqcap B \equiv^? B$ that are generated by an application of the decomposition rule are called *D-equations*. Equations that are elements of the input problem Γ_0 are called *I-equations*. Any equation e' that is not an I-equation or a D-equation has a unique predecessor w.r.t. $\to$, i.e., there is an equation $e \in \mathcal{E}$ such that $e \to e'$.

Starting with the set $\mathcal{F} := \widehat{\Gamma}$ we will now step by step extend $\mathcal{F}$ by a predecessor of an equation in $\mathcal{F}$ until no new predecessors can be added. Since $\mathcal{E}$ is finite, this process terminates after a finite number of steps. After termination we have $\mathcal{E} = \mathcal{F}$, and thus in particular $\Gamma_0 \subseteq \mathcal{F}$. This is due to the fact that, for every element e_0 of $\mathcal{E}$, there are $n \geq 0$ elements $e_1, \ldots, e_n \in \mathcal{E}$ such that $e_0 \to e_1 \to \ldots \to e_n$ and $e_n \in \widehat{\Gamma}$. Thus, it is enough to show that the set $\mathcal{F}$ satisfies the following *invariant*:

$$(*) \quad \text{the substitution } \sigma \text{ solves every equation in } \mathcal{F}.$$

Since σ is a solution of $\widehat{\Gamma}$, this invariant is initially satisfied. To prove that it is preserved under adding predecessors of equations in $\mathcal{F}$, we start with the equations of minimal role depth. To be more precise, if the equation e is of the form $C \equiv^? D$, we define the *role depth of e w.r.t. σ* to be the role depth[10] of the concept term $\sigma(C) \sqcap \sigma(D)$. The strict order $\succ$ on $\mathcal{E}$ is defined as follows: $e \succ_\sigma e'$ iff the role depth of e w.r.t. σ is larger than the role depth

---

[10] see the proof of Proposition 3.5 for a definition.

of e' w.r.t. σ. We write $e \approx_\sigma e'$ if e and e' have the same role depth w.r.t. σ. The following is an easy consequence of the definition of σ and of our rules:

$$(**) \quad e_1 \to e_2 \to \ldots \to e_n \quad \text{implies} \quad e_1 \approx_\sigma e_2 \approx_\sigma \ldots \approx_\sigma e_n.$$

Assume that we have already constructed a set $\mathcal{F}$ such that the invariant $(*)$ is satisfied. Let e' be an equation in $\mathcal{F}$ such that

- there is an $e \in \mathcal{E} \setminus \mathcal{F}$ with $e \to e'$;
- e' is of minimal role depth with this property, i.e., if $f' \in \mathcal{F}$ is such that $e' \succ f'$ and f' has a predecessor f w.r.t. $\to$, then $f \in \mathcal{F}$.

If no such equation e' exists, then we are finished, and we have $\mathcal{E} = \mathcal{F}$. Otherwise, let e' be such an equation and e its predecessor w.r.t. $\to$. We add e to $\mathcal{F}$. In order to show that the invariant $(*)$ is still satisfied, we make a case distinction according to which rule was applied to e to produce e':

(1) *Eager-Assignment.* By an application of this rule, the assignment for X is modified from $S_X = \emptyset$ to $S_X = \{A_1, \ldots, A_n\}$, where $A_1, \ldots, A_n$ are non-variable atoms. In addition, X is labeled as finished. Since the assignment of a finished variable cannot be changed anymore, we know that we also have $S_X = \{A_1, \ldots, A_n\}$ in the final assignment, and thus $\sigma(X) = \sigma(A_1) \sqcap \ldots \sqcap \sigma(A_n)$. The rule modifies equations as follows: all equations containing X are expanded w.r.t. the assignment $S_X = \{A_1, \ldots, A_n\}$. Since e is transformed into e' using this rule, it must contain X. We assume for the sake of simplicity that X is contained in the left-hand side of e, but not in the right-hand side, i.e., e is of the form $C \sqcap X \equiv^? D$ and the new equation $e' \in \Gamma'$ obtained from e is $C \sqcap X \sqcap A_1 \sqcap \ldots \sqcap A_n \equiv^? D$. Since σ solves e', we have $\sigma(D) \equiv \sigma(C \sqcap X \sqcap A_1 \sqcap \ldots \sqcap A_n) \equiv \sigma(C) \sqcap \sigma(A_1) \sqcap \ldots \sqcap \sigma(A_n) \sqcap \sigma(A_1) \sqcap \ldots \sqcap \sigma(A_n) \equiv \sigma(C) \sqcap \sigma(A_1) \sqcap \ldots \sqcap \sigma(A_n) \equiv \sigma(C \sqcap X)$, which shows that σ also solves e.

(2) *Decomposition.* Without loss of generality, we consider the L-variant of this rule. Thus, the equation e is of the form $D \sqcap \exists r.C \equiv^? E \sqcap \exists r.B$, and e' is obtained from e by adding $\exists r.C$ to the right-hand side, i.e., e' is of the form $D \sqcap \exists r.C \equiv^? E \sqcap \exists r.B \sqcap \exists r.C$. We know that σ solves e'. Thus, if we can show $\sigma(B) \sqsubseteq \sigma(C)$, then we have $\sigma(D \sqcap \exists r.C) \equiv \sigma(E) \sqcap \sigma(\exists r.B) \sqcap \sigma(\exists r.C) \equiv \sigma(E) \sqcap \sigma(\exists r.B)$, which shows that σ solves e.

Consequently, it is sufficient to prove $\sigma(B) \sqsubseteq \sigma(C)$. The Decomposition rule also generates the equation $C \sqcap B \equiv^? B$ and expands it w.r.t. the assignments of all the variables contained in this equation, unless this equation has already been generated before. Thus, either this application or a previous one of the Decomposition rule has generated the equation $C \sqcap B \equiv^? B$, and then expanded it (w.r.t. the current assignment at that time) to an equation e_1. Since atoms are never removed from an assignment, the atoms present in the assignment at the time when the Decomposition rule generated the equation $C \sqcap B \equiv^? B$ are also present in the final assignment used to define the substitution σ. Thus, if we can show that σ solves e_1, then we have also shown that σ solves $C \sqcap B \equiv^? B$, and thus satisfies $\sigma(B) \sqsubseteq \sigma(C)$.

Since equations are never completely removed by our rules, but only modified, there is a sequence of equations $e_1 \to e_2 \to \ldots \to e_n$ such that $e_n \in \widehat{\Gamma}$. Property $(**)$ thus yields $e_1 \approx_\sigma e_2 \approx_\sigma \ldots \approx_\sigma e_n$. In addition, the role depth of $C \sqcap B \equiv^? B$ w.r.t. σ is the same as the role depth of e_1 w.r.t. σ. Consequently, we have $e' \succ e_i$ for all $i, 1 \le i \le n$. Now, assume that $e_1 \notin \mathcal{F}$. Then there is an $i > 1$ such that $e_i \in \mathcal{F}$, but $e_{i-1} \in \mathcal{E} \setminus \mathcal{F}$.

This contradicts our assumption that e' is minimal. Thus, we have shown that $e_1 \in \mathcal{F}$, and this implies that σ solves e_1.

Overall, this finishes the proof that σ solves e.

(3) *Extension.* By an application of this rule, the assignment for X is modified by adding a non-variable atom A to it. Since atoms are never removed from an assignment, we know that we also have $A \in S_X$ in the final assignment, and thus $\sigma(X) \sqsubseteq \sigma(A)$. The rule modifies equations as follows: all equations containing X are expanded w.r.t. the new assignment for X. Since e is transformed into e' using this rule, it must contain X. We assume for the sake of simplicity that X is contained in the left-hand side of e, but not in the right-hand side, i.e., e is of the form $C \sqcap X \equiv^? D$ and the new equation e' obtained from e is $C \sqcap X \sqcap A \equiv^? D$. Since σ solves e', we have $\sigma(D) \equiv \sigma(C \sqcap X \sqcap A) \equiv \sigma(C) \sqcap \sigma(X) \sqcap \sigma(A) \equiv \sigma(C) \sqcap \sigma(X) \equiv \sigma(C \sqcap X)$, which shows that σ also solves e.

To sum up, we have shown that the invariant $(*)$ is still satisfied after adding e to $\mathcal{F}$. This completes the proof of soundness of our procedure. $\qquad \square$

It remains to show completeness of Algorithm 6.2. Thus, assume that the input unification problem Γ_0 is solvable. Proposition 5.4 tells us that Γ_0 then has an is-minimal reduced ground unifier γ, and Proposition 5.11 implies that, for every variable X occurring in Γ_0, there is a set S_X^γ of non-variable atoms of Γ_0 such that

$$\gamma(X) \equiv \gamma(\bigsqcap S_X^\gamma),$$

where, for a set of non-variable atoms S of Γ_0, the expression $\bigsqcap S$ denotes the conjunction of the elements of S (where the empty conjunction is $\top$).

Lemma 6.6 (Completeness)**.** *Let Γ_0 be a flat $\mathcal{EL}$-unification problem, and assume that γ is an is-minimal reduced ground unifier of Γ_0. Then there is a successful run of Algorithm 6.2 on input Γ_0 that returns a unifier σ that is equivalent to γ, i.e., satisfies $\sigma(X) \equiv \gamma(X)$ for all variables X occurring in Γ_0.*

Proof. The algorithm starts with $\Gamma := \Gamma_0$ and the initial assignment $S_X := \emptyset$ for all variables X occurring in Γ_0. It then applies rules that change Γ and the current assignment as long as the problem Γ contains an unsolved equation.

We use γ to guide the (don't know) non-deterministic choices to be made during the algorithm. We show that this ensures that the run of the algorithm generated this way *does not fail* and that the following *invariants* are satisfied throughout this run:

(I_1) γ is a unifier of Γ;
(I_2) for all atoms $B \in S_X$ there exists an atom $A \in S_X^\gamma$ such that $\gamma(A) \sqsubseteq \gamma(B)$;
(I_3) for all finished variables X we have $\gamma(X) \equiv \gamma(\bigsqcap S_X)$.

Before constructing a run that satisfies these invariants, let us point out two interesting consequences that they have:

(C_1) *The current assignment is always acyclic.* In fact, if X directly depends on Y, then there is an atom $B \in S_X$ that has the form $B = \exists r.Y$ for some role name r. Invariant I_2 then implies that there is an $A \in S_X^\gamma$ such that $\gamma(X) \sqsubseteq \gamma(A) \sqsubseteq \gamma(B) = \exists r.\gamma(Y)$. Thus, if X depends on X, then there are $k \geq 1$ role names $r_1, \ldots, r_k$ such that $\gamma(X) \sqsubseteq \exists r_1. \cdots \exists r_k.\gamma(X)$, which is impossible.

(C_2) *For each variable X occurring in Γ_0, we have $\gamma(X) \sqsubseteq \sigma(X)$, where σ is the current substitution induced by the current assignment.* This is again a consequence of invariant I_2. Indeed, recall that the fact that the current assignment is acyclic implies that there is a strict linear order $>$ on the variables occurring in Γ such that $X > Y$ if X depends on Y. The current substitution σ is defined along this order. We prove $\gamma(X) \sqsubseteq \sigma(X)$ by induction on this order.

Consider the least variable X. If $S_X = \emptyset$, then $\sigma(X) = \top$, and thus $\gamma(X) \sqsubseteq \sigma(X)$ is trivially satisfied. Otherwise, we know, for every $B \in S_X$, that it does not contain any variables, which implies that $\sigma(B) = B = \gamma(B) \sqsupseteq \gamma(A)$ for some atom $A \in S_X^\gamma$. Obviously, this yields $\sigma(X) = \sigma(\sqcap S_X) \sqsupseteq \gamma(\sqcap S_X^\gamma) = \gamma(X)$.

Now, assume that $\gamma(Y) \sqsubseteq \sigma(Y)$ holds for all variables $Y < X$. Since the concept constructors of $\mathcal{EL}$ are monotone w.r.t. subsumption, this implies $\gamma(C) \sqsubseteq \sigma(C)$ for all concept terms C containing only variables smaller than X. If S_X is empty, then $\sigma(X) = \top \sqsupseteq \gamma(X)$ is trivially satisfied. Otherwise, we know, for every $B \in S_X$, that it contains only variables smaller than X. This yields $\sigma(B) \sqsupseteq \gamma(B) \sqsupseteq \gamma(A)$ for some atom $A \in S_X^\gamma$. Again, this implies $\sigma(X) = \sigma(\sqcap S_X) \sqsupseteq \gamma(\sqcap S_X^\gamma) = \gamma(X)$.

Since γ was assumed to be an is-minimal unifier of Γ_0, the consequence C_2 implies that σ can only be a unifier of Γ_0 if σ is equivalent to γ. If the run has terminated successfully, then the final substitution σ obtained by the run is a unifier of Γ_0 (due to soundness). Thus, in this case the computed unifier σ is indeed equivalent to γ. Consequently, to prove the lemma, it is sufficient to *construct a non-failing run of the algorithm that satisfies the above invariants.*

The invariants are initially satisfied since γ is a unifier of Γ_0, the initial assignment for all variables X occurring in Γ_0 is $S_X = \emptyset$, and there are no finished variables. Now, assume that, by application of the rules of Algorithm 6.2, we have constructed a unification problem Γ and a current assignment such that the invariants are satisfied.

(1) If all equations in Γ are solved, then the run terminates successfully, and we are done.

(2) If there is an unsolved equation to which the *Eager-Assignment rule* applies, then the algorithm picks such an equation e and applies this rule to it. Without loss of generality, we assume that the L-variant of the rule is applied. The selected equation e is of the form

$$X \sqcap Z_1 \sqcap \ldots \sqcap Z_k \equiv^? A_1 \sqcap \ldots \sqcap A_n \sqcap Y_1 \sqcap \ldots \sqcap Y_m,$$

where $A_1, \ldots, A_n$ are non-variable atoms, and $Y_1, \ldots, Y_m, Z_1, \ldots Z_k$ are finished variables. Because the left-hand side of the equation does not contain any non-variable atoms, we know that $S_X = S_{Z_1} = \ldots = S_{Z_k} = \emptyset$ (since the algorithm keeps all equations expanded). Since $Z_1, \ldots, Z_k$ are finished, we thus have $\gamma(Z_1) = \ldots = \gamma(Z_k) = \top$ (by invariant I_3). We also know that $S_{Y_i} \subseteq \{A_1, \ldots, A_n\}$ for all $i, 1 \le i \le m$. Since the variables Y_i are finished, invariant I_3 implies that $\gamma(Y_i) \sqsupseteq \gamma(A_1) \sqcap \ldots \sqcap \gamma(A_n)$.

The new assignment for X is $S_X = \{A_1, \ldots, A_n\}$, all equations containing X are expanded w.r.t. this assignment, and X becomes a finished variable. First, we show that I_3 is satisfied. Nothing has changed for the variables that were already finished before the application of the rule. However, X is now also finished. Thus, we must show that $\gamma(X) \equiv \gamma(\sqcap S_X)$. We know that γ solves the equation e (by I_1). This yields $\gamma(X) \equiv \gamma(X) \sqcap \gamma(Z_1) \sqcap \ldots \sqcap \gamma(Z_k) \equiv \gamma(A_1) \sqcap \ldots \sqcap \gamma(A_n) \sqcap \gamma(Y_1) \sqcap \ldots \sqcap \gamma(Y_m) \equiv \gamma(A_1) \sqcap \ldots \sqcap \gamma(A_n) = \gamma(\sqcap S_X)$. Regarding I_2, the only assignment that was changed is the one for X. Since the new assignment for X is $S_X = \{A_1, \ldots, A_n\}$, and we have

already shown that $\gamma(X) \equiv \gamma(A_1) \sqcap \ldots \sqcap \gamma(A_n)$, the invariant I_2 holds by Corollary 3.2. Note that this also implies that the new assignment is acyclic, and thus the application of the Eager-Assignment rule does not fail. Finally, consider the invariant I_1. The rule application modifies equations containing X by adding the atoms $A_1, \ldots, A_n$. Since $\gamma(X) \equiv \gamma(A_1) \sqcap \ldots \sqcap \gamma(A_n)$, an equation that was solved by γ before this modification, is also solved by γ after this modification. To sum up, we have shown that the application of the Eager-Assignment rule does not fail and preserves the invariants.

(3) If there is *no* unsolved equation to which the *Eager-Assignment rule* applies, then the algorithm picks an unsolved equation e and an unsolved atom A occurring in this equation. We must show that we can apply either the Decomposition or the Extension rule to A in e such that the invariants stay satisfied. Without loss of generality, we assume that the unsolved atom A occurs on the left-hand side of the equation e.

(a) First, assume that A *is an existential restriction* $A = \exists r.C$. The selected unsolved equation e is thus of the form

$$\exists r.C \sqcap A_1 \sqcap \ldots \sqcap A_m \equiv B_1 \sqcap \ldots \sqcap B_n,$$

where $A_1, \ldots, A_m$ and $B_1, \ldots, B_n$ are (variable or non-variable) atoms and $\exists r.C \notin \{B_1, \ldots, B_n\}$. Since γ solves this equation (by invariant I_1), Corollary 3.2 implies that there must be an $i, 1 \le i \le n$, such that $\gamma(B_i) \sqsubseteq \exists r.\gamma(C)$.

 (i) If B_i is an existential restriction $B_i = \exists r.B$, then we have $\gamma(B) \sqsubseteq \gamma(C)$. We apply the Decomposition rule to A and B_i. The application of this rule modifies the equation e to an equation e' by adding the atom A to the right-hand side. In addition, it generates the equation $C \sqcap B \equiv^? B$ and expands it w.r.t. the assignments of all variables contained in this equation (unless this equation has been generated before). After the application of this rule, the invariants I_2 and I_3 are still satisfied since the current assignments and the set of finished variables remain unchanged. Regarding invariant I_1, since γ solves e, it obviously also solves e' due to the fact that $\gamma(B_i) \sqsubseteq \gamma(A)$ and B_i is a conjunct on the right-hand side of e. In addition, $\gamma(B) \sqsubseteq \gamma(C)$ implies that γ also solves the equation $C \sqcap B \equiv^? B$. Since invariant I_2 is satisfied, this implies that γ also solves the equation obtained from $C \sqcap B \equiv^? B$ by expanding it w.r.t. the assignments of all variables contained in it.

 (ii) Assume that there is no $i, 1 \le i \le n$, such that B_i is an existential restriction satisfying $\gamma(B_i) \sqsubseteq \gamma(A)$. Thus, if B_i is such that $\gamma(B_i) \sqsubseteq \gamma(A)$, then we know that $B_i = X$ is a variable. We want to apply the Extension rule to A and X. To be able to do this, we must first show that X is not a finished variable. Thus, assume that X is finished, and let $S_X = \{C_1, \ldots, C_\ell\}$. Invariant I_3 yields $\gamma(C_1) \sqcap \ldots \sqcap \gamma(C_\ell) = \gamma(X) = \gamma(B_i) \sqsubseteq \gamma(A) = \exists r.\gamma(C)$, and thus there is a $j, 1 \le j \le \ell$, such that $\gamma(C_j) \sqsubseteq \gamma(A)$. Since A is an existential restriction, the non-variable atom C_j must also be an existential restriction, and since the equation e is expanded, $C_j \in S_X$ occurs on the right-hand side of this equation. This contradicts our assumption that there is no such existential restriction on the right-hand side. Thus, we have shown that X is not finished, which means that we can apply the Extension rule to A and X.

 The application of this rule adds the atom A to the assignment for X, and it expands all equations containing X w.r.t. this new assignment, i.e., it adds A to the left-hand side and/or right-hand side of an equation whenever X is

contained in this side. Since we know that $\gamma(X) \sqsubseteq \gamma(A)$, it is easy to see that, if γ solves an equation before this expansion, it also solves it after the expansion. Thus invariant I_1 is satisfied. Invariant I_2 also remains satisfied. In fact, if $S_X^\gamma = \{D_1, \ldots, D_k\}$, then $\gamma(D_1) \sqcap \ldots \sqcap \gamma(D_k) = \gamma(X) \sqsubseteq \gamma(A)$ implies that there is a $j, 1 \le j \le \ell$, such that $\gamma(D_j) \sqsubseteq \gamma(A)$. The fact that I_2 is satisfied by the new assignment also implies that this new assignment is acyclic, and thus the application of the Extension rule does not fail. Invariant I_3 is still satisfied since X is not finished, and the assignments of variables different from X were not changed.

(b) Second, assume that *A is a concept name.* The selected unsolved equation e is thus of the form

$$A \sqcap A_1 \sqcap \ldots \sqcap A_m \equiv B_1 \sqcap \ldots \sqcap B_n,$$

where $A_1, \ldots, A_m$ and $B_1, \ldots, B_n$ are (variable or non-variable) atoms, and $A \notin \{B_1, \ldots, B_n\}$. Since γ solves this equation (by invariant I_1), Corollary 3.2 implies that there must be an $i, 1 \le i \le n$, such that $\gamma(B_i) \sqsubseteq \gamma(A) = A$. Since $A \notin \{B_1, \ldots, B_n\}$, we know that $B_i = X$ is a variable. We want to apply the Extension rule to A and X. To be able to do this, we must first show that X is not a finished variable.

Thus, assume that X is finished, and let $S_X = \{C_1, \ldots, C_\ell\}$. Invariant I_3 yields $\gamma(C_1) \sqcap \ldots \sqcap \gamma(C_\ell) = \gamma(X) = \gamma(B_i) \sqsubseteq \gamma(A) = A$, and thus there is a $j, 1 \le j \le \ell$, such that $\gamma(C_j) \sqsubseteq A$. Since A is a concept name, the non-variable atom C_j must actually be equal to A, and since the equation e is expanded, $C_j = A \in S_X$ occurs on the right-hand side of this equation. This contradicts our assumption that A is an unsolved atom. Thus, we have shown that X is not finished, which means that we can apply the Extension rule to A and X. The application of this rule adds the atom A to the assignment for X, and it expands all equations containing X w.r.t. this new assignment. The proof that this rule application does not fail and preserves the invariants is identical to the one for the case where A was an existential restriction. ☐

To sum up, we have shown that Algorithm 6.2 always terminates (in non-deterministic polynomial time) and that it is sound and complete. This finishes the proof of Theorem 6.3.

7. Unification in semilattices with monotone operators

Unification problems and their types were originally not introduced for Description Logics, but for equational theories [16]. In this section, we show that the above results for unification in $\mathcal{EL}$ can actually be viewed as results for an equational theory. As shown in [28], the equivalence problem for $\mathcal{EL}$-concept terms corresponds to the word problem for the equational theory of semilattices with monotone operators. In order to define this theory, we consider a signature Σ_{SLmO} consisting of a binary function symbol $\wedge$, a constant symbol 1, and finitely many unary function symbols $f_1, \ldots, f_n$. Terms can then be built using these symbols and additional variable symbols and free constant symbols.

Definition 7.1. The equational theory of *semilattices with monotone operators* is defined by the following identities:

$$SLmO := \{x \wedge (y \wedge z) = (x \wedge y) \wedge z, \; x \wedge y = y \wedge x, \; x \wedge x = x, \; x \wedge 1 = x\} \cup$$

$$\{f_i(x \wedge y) \wedge f_i(y) = f_i(x \wedge y) \mid 1 \leq i \leq n\}$$

A given $\mathcal{EL}$-concept term C using only roles $r_1, \ldots, r_n$ can be translated into a term t_C over the signature Σ_{SLmO} by replacing each concept constant A by a corresponding free constant a, each concept variable X by a corresponding variable x, $\top$ by 1, $\sqcap$ by $\wedge$, and $\exists r_i$ by f_i. For example, the $\mathcal{EL}$-concept term $C = A \sqcap \exists r_1.\top \sqcap \exists r_3.(X \sqcap B)$ is translated into $t_C = a \wedge f_1(1) \wedge f_3(x \wedge b)$. Conversely, any term over the signature Σ_{SLmO} can be translated back into an $\mathcal{EL}$-concept term.

Lemma 7.2. *Let C, D be $\mathcal{EL}$-concept term using only roles $r_1, \ldots, r_n$. Then $C \equiv D$ iff $t_C =_{SLmO} t_D$.*

As an immediate consequence of this lemma, we have that unification in the DL $\mathcal{EL}$ corresponds to unification modulo the equational theory $SLmO$. Thus, Theorem 4.1 implies that $SLmO$ has unification type zero, and Theorem 5.14 implies that $SLmO$-unification is NP-complete.

Corollary 7.3. *The equational theory $SLmO$ of semilattices with monotone operators has unification type zero, and deciding solvability of an $SLmO$-unification problem is an NP-complete problem.*

Since the unification problem introduced in Theorem 4.1 contains only one role r, this is already true in the presence of a single monotone operator.

8. Conclusion

In this paper, we have shown that unification in the DL $\mathcal{EL}$ is of type zero and NP-complete. There are interesting differences between the behavior of $\mathcal{EL}$ and the closely related DL $\mathcal{FL}_0$ w.r.t. unification and matching. Though the unification types coincide for these two DLs, the complexities of the decision problems differ: $\mathcal{FL}_0$-unification is ExpTime-complete, and thus considerably harder than $\mathcal{EL}$-unification. In contrast, $\mathcal{FL}_0$-matching is polynomial, and thus considerably easier than $\mathcal{EL}$-matching, which is NP-complete. In addition to showing the complexity upper bound for $\mathcal{EL}$-unification by a simple "guess and then test" NP-algorithm, we have also developed a more goal-oriented NP-algorithm that makes (don't know) non-deterministic decisions (i.e., ones that require backtracking) only if they are triggered by unsolved atoms in the unification problem.

As future work, we will consider also unification of concept terms for other members of the $\mathcal{EL}$-family of DLs [5]. In addition, we will investigate unification modulo more expressive terminological formalisms. On the practical side, we will optimize and implement the goal-oriented $\mathcal{EL}$-unification algorithm developed in Section 6. We intend to test the usefulness of this algorithm for the purpose on finding redundancies in $\mathcal{EL}$-based ontologies by considering extensions of the medical ontology SNOMED CT. For example, in [18], two different extensions of SNOMED CT by so-called post-coordinated concepts were considered. The authors used an (incomplete) equivalence test to find out how large the overlap between the two extensions is (i.e., how many of the new concepts belonged to both extensions). As pointed out in the introduction, the equivalence test cannot deal with situations where

different knowledge engineers use different names for concepts, or model on different levels of granularity. We want to find out whether using unifiability rather than equivalence finds more cases of overlapping concepts. Of course, in the case of unification one may also obtain false positives, i.e., pairs of concepts that are unifiable, but are not meant to represent the same (intuitive) concept. It is also important to find out how often this happens. Another problem to be dealt with in this application is the development of heuristics for choosing the pairs of concepts to be tested for unifiability and for deciding which concept names are turned into variables.

References

[1] Franz Baader. Characterizations of unification type zero. In N. Dershowitz, editor, *Proceedings of the 3rd International Conference on Rewriting Techniques and Applications*, volume 355 of *Lecture Notes in Computer Science*, pages 2–14, Chapel Hill, North Carolina, 1989. Springer-Verlag.

[2] Franz Baader. Unification in commutative theories. *J. Symbolic Computation*, 8(5):479–497, 1989.

[3] Franz Baader. Terminological cycles in KL-ONE-based knowledge representation languages. In *Proc. of the 8th Nat. Conf. on Artificial Intelligence (AAAI'90)*, pages 621–626, Boston (Ma, USA), 1990.

[4] Franz Baader. Terminological cycles in a description logic with existential restrictions. In Georg Gottlob and Toby Walsh, editors, *Proc. of the 18th Int. Joint Conf. on Artificial Intelligence (IJCAI 2003)*, pages 325–330, Acapulco, Mexico, 2003. Morgan Kaufmann, Los Altos.

[5] Franz Baader, Sebastian Brandt, and Carsten Lutz. Pushing the $\mathcal{EL}$ envelope. In Leslie Pack Kaelbling and Alessandro Saffiotti, editors, *Proc. of the 19th Int. Joint Conf. on Artificial Intelligence (IJCAI 2005)*, pages 364–369, Edinburgh (UK), 2005. Morgan Kaufmann, Los Altos.

[6] Franz Baader, Diego Calvanese, Deborah McGuinness, Daniele Nardi, and Peter F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press, 2003.

[7] Franz Baader and Ralf Küsters. Matching in description logics with existential restrictions. In *Proc. of the 7th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR 2000)*, pages 261–272, 2000.

[8] Franz Baader and Ralf Küsters. Unification in a description logic with transitive closure of roles. In Robert Nieuwenhuis and Andrei Voronkov, editors, *Proc. of the 8th Int. Conf. on Logic for Programming and Automated Reasoning (LPAR 2001)*, volume 2250 of *Lecture Notes in Artificial Intelligence*, pages 217–232, Havana, Cuba, 2001. Springer-Verlag.

[9] Franz Baader, Ralf Küsters, Alex Borgida, and Deborah L. McGuinness. Matching in description logics. *J. of Logic and Computation*, 9(3):411–447, 1999.

[10] Franz Baader, Ralf Küsters, and Ralf Molitor. Computing least common subsumers in description logics with existential restrictions. In *Proc. of the 16th Int. Joint Conf. on Artificial Intelligence (IJCAI'99)*, pages 96–101, 1999.

[11] Franz Baader and Barbara Morawska. Unification in the description logic $\mathcal{EL}$. In Ralf Treinen, editor, *Proc. of the 20th Int. Conf. on Rewriting Techniques and Applications (RTA 2009)*, volume 5595 of *Lecture Notes in Computer Science*, pages 350–364. Springer-Verlag, 2009.

[12] Franz Baader and Paliath Narendran. Unification of concepts terms in description logics. *J. of Symbolic Computation*, 31(3):277–305, 2001.

[13] Franz Baader and Tobias Nipkow. *Term Rewriting and All That*. Cambridge University Press, United Kingdom, 1998.

[14] Franz Baader and Werner Nutt. Basic description logics. In [6], pages 43–95. 2003.

[15] Franz Baader, Baris Sertkaya, and Anni-Yasmin Turhan. Computing the least common subsumer w.r.t. a background terminology. *J. of Applied Logic*, 5(3):392–420, 2007.

[16] Franz Baader and Wayne Snyder. Unification theory. In J.A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, volume I, pages 447–533. Elsevier Science Publishers, 2001.

[17] Sebastian Brandt. Polynomial time reasoning in a description logic with existential restrictions, GCI axioms, and—what else? In Ramon López de Mántaras and Lorenza Saitta, editors, *Proc. of the 16th Eur. Conf. on Artificial Intelligence (ECAI 2004)*, pages 298–302, 2004.

[18] James R. Campbell, Alejandro Lopez Osornio, Fernan de Quiros, Daniel Luna, and Guillermo Reynoso. Semantic interoperability and SNOMED CT: A case study in clinical problem lists. In K.A. Kuhn, J.R. Warren, and T.-Y. Leong, editors, *Proc. of the 12th World Congress on Health (Medical) Informatics (MEDINFO 2007)*, pages 2401–2402. IOS Press, 2007.

[19] Silvio Ghilardi. Best solving modal equations. *Ann. Pure Appl. Logic*, 102(3):183–198, 2000.

[20] Ian Horrocks, Peter F. Patel-Schneider, and Frank van Harmelen. From SHIQ and RDF to OWL: The making of a web ontology language. *Journal of Web Semantics*, 1(1):7–26, 2003.

[21] Ian Horrocks, Ulrike Sattler, and Stefan Tobies. Practical reasoning for very expressive description logics. *J. of the Interest Group in Pure and Applied Logic*, 8(3):239–264, 2000.

[22] Jean-Pierre Jouannaud and Claude Kirchner. Solving equations in abstract algebras: A rule-based survey of unification. In J.-L. Lassez and G. Plotkin, editors, *Computational Logic: Essays in Honor of A. Robinson*. MIT Press, Cambridge, MA, 1991.

[23] Yevgeny Kazakov and Hans de Nivelle. Subsumption of concepts in $\mathcal{FL}_0$ for (cyclic) terminologies with respect to descriptive semantics is PSPACE-complete. In *Proc. of the 2003 Description Logic Workshop (DL 2003)*. CEUR Electronic Workshop Proceedings, http://CEUR-WS.org/Vol-81/, 2003.

[24] Ralf Küsters. *Non-standard Inferences in Description Logics*, volume 2100 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, 2001.

[25] Hector J. Levesque and Ron J. Brachman. A fundamental tradeoff in knowledge representation and reasoning. In Ron J. Brachman and Hector J. Levesque, editors, *Readings in Knowledge Representation*, pages 41–70. Morgan Kaufmann, Los Altos, 1985.

[26] Bernhard Nebel. Terminological reasoning is inherently intractable. *Artificial Intelligence*, 43:235–249, 1990.

[27] Alan Rector and Ian Horrocks. Experience building a large, re-usable medical ontology using a description logic with transitivity and concept inclusions. In *Proceedings of the Workshop on Ontological Engineering, AAAI Spring Symposium (AAAI'97)*, Stanford, CA, 1997. AAAI Press.

[28] Viorica Sofronie-Stokkermans. Locality and subsumption testing in $\mathcal{EL}$ and some of its extensions. In *Proc. Advances in Modal Logic (AiML'08)*, 2008.

[29] Frank Wolter and Michael Zakharyaschev. Undecidability of the unification and admissibility problems for modal and description logics. *ACM Trans. Comput. Log.*, 9(4), 2008.