# Simplifying Description Logic Ontologies

Nadeschda Nikitina[1] and Sven Schewe[2]

[1] University of Oxford, UK
[2] University of Liverpool, UK

**Abstract.** We discuss the problem of minimizing TBoxes expressed in the lightweight description logic $\mathcal{EL}$, which forms a basis of some large ontologies like SNOMED, Gene Ontology, NCI and Galen. We show that the minimization of TBoxes is intractable (NP-complete). While this looks like a bad news result, we also provide a heuristic technique for minimizing TBoxes. We prove the correctness of the heuristics and show that it provides optimal results for a class of ontologies, which we define through an acyclicity constraint over a reference relation between equivalence classes of concepts. To establish the feasibility of our approach, we have implemented the algorithm and evaluated its effectiveness on a small suite of benchmarks.

## 1 Introduction

It is well-known that the same facts can be represented in many different ways, and that the size of these representations can vary significantly. This is also reflected in ontology engineering, where the syntactic form of ontologies can be more complex than necessary. For instance, throughout the development (and the life-cycle) of an ontology, the way in which concepts and the relationship between them are represented within the ontology are constantly changing. For example, a name for a complex concept expression is often introduced only after it has been used several times and has proved to be important. Another example are dependencies between concepts that evolve over time, resulting in new subsumption relations between concepts ($A_1 \sqsubseteq A_2$). As a result, previously reasonable concept expressions may become unnecessarily complex. In the given example, $A_1 \sqcap A_2$ becomes equivalent to $A_1$.

Clearly, unnecessary complexity impacts on the maintenance effort as well as the usability of ontologies. For instance, keeping track of dependencies between complex concept expressions and relationships between them is more cumbersome when it contains unnecessarily complex or unnecessarily many different concept expressions. As a result, the chance of introducing unwanted consequences is higher. Moreover, unintended redundancy decreases the overall quality of the ontology.

Removing unnecessary syntactic complexity from ontologies by hand is a difficult task: for the average ontology, it is almost impossible to obtain the minimal representation without tool support. Thus, automated methods that help to assess the current succinctness of an ontology and generate suggestions on how to increase it would be highly valued by ontology engineers.

It is easy to envision scenarios that demonstrate the usefulness of rewriting for reducing the cognitive complexity of axioms. For instance, when a complex concept $C$ is frequently used in the axioms of an ontology and there is an equivalent atomic concept $A_C$, the ontology will diminish in size when occurrences of $C$ are replaced by $A_C$.

*Example 1.* Consider the following excerpt from the ontology Galen [1]:

$$\begin{aligned}\texttt{Clotting} \sqsubseteq\ & \exists\,\texttt{actsSpecificallyOn}.(\texttt{Blood} \sqcap \exists\,\texttt{hasPhysicalState}. \\ & \qquad (\texttt{PhysicalState} \sqcap \exists\,\texttt{hasState.liquid})) \sqcap \\ & \exists\texttt{hasOutcome}.(\texttt{Blood} \sqcap \exists\,\texttt{hasPhysicalState.solidState})\end{aligned} \quad (1)$$

$$\texttt{LiquidState} \equiv \texttt{PhysicalState} \sqcap \exists\,\texttt{hasState.liquid} \quad (2)$$

$$\texttt{LiquidBlood} \equiv \texttt{Blood} \sqcap \exists\,\texttt{hasPhysicalState.LiquidState} \quad (3)$$

Given concepts defined in Axioms 2 and 3 above, we can easily rewrite Axiom 1 to obtain the following, simpler axiom containing only 6 references to concepts and roles (as opposed to 10 references in Axiom 1):

$$\begin{aligned}\texttt{Clotting} \sqsubseteq\ & \exists\,\texttt{actsSpecificallyOn.LiquidBlood} \sqcap \\ & \exists\texttt{hasOutcome}.(\texttt{Blood} \sqcap \exists\,\texttt{hasPhysicalState.solidState})\end{aligned} \quad (4)$$

In description logics [2], few results towards simplifying ontologies have been obtained so far. Grimm et al. [3] propose an algorithm for eliminating semantically redundant axioms from ontologies. In the above approach, axioms are considered as atoms that cannot be split into parts or changed in any other way. With the specific goal of improving reasoning efficiency, Bienvenu et al. [4] propose a normal form called prime implicates normal form for $\mathcal{ALC}$ ontologies. However, as a side-effect of this transformation, a doubly-exponential blowup in concept size can occur.

In this paper, we investigate the succinctness for the lightweight description logic $\mathcal{EL}$. The tractable OWL 2 EL profile [5] of the W3C-specified OWL Web Ontology Language [6] is based on DLs of the $\mathcal{EL}$ family [7]. We consider the problem of finding a minimal equivalent representation for a given $\mathcal{EL}$ ontology. First, we demonstrate that we can reduce the size of a representation by up to an exponent even in the case that the ontology does not contain any redundant axioms. We show that the related decision problem (is there an equivalent ontology of size $\leq k$?) is NP-complete by a reduction from the set cover problem, which is one of the standard NP-complete problems. We also show that, just as for other reasoning problems in $\mathcal{EL}$, ontology minimization becomes simpler under the absence of a particular type of cycles. We identify a class of TBoxes, for which the problem can be solved in PTIME instead of NP and implement a tractable algorithm that computes a minimal TBox for this class of TBoxes. The algorithm can also be applied to more expressive and most cyclic TBoxes[3], however without a guarantee of minimality. We apply an implementation of the algorithm to various existing ontologies and show that their succinctness can be improved. For instance, in case of Galen, we managed to reduce the number of complex concepts occurrences by 955 and the number of references to atomic concepts and roles by 1130.

[3] The extension to general TBoxes is a trivial modification of the algorithm

The paper is organized as follows: In Section 2, we recall the necessary preliminaries on description logics. Section 3 demonstrates the potential of minimization. In the same section, we also introduce the basic definitions of the size of ontologies and formally state the corresponding decision problem. In Section 4, we derive the complexity bounds for this decision problem. Section 5 defines the class of TBoxes, for which the problem can be solved in PTIME instead of NP and presents a tractable algorithm that computes a minimal TBox for this class of TBoxes. In Section 6, we present experimental results for a selection of ontologies. Finally, we discus related approaches in Section 7 before we conclude and outline future work in Section 8.

## 2 Preliminaries

We recall the basic notions in description logics [2] required in this paper. Let $N_C$ and $N_R$ be countably infinite and mutually disjoint sets of concept symbols and role symbols. An $\mathcal{EL}$ concept $C$ is defined as

$$C ::= A|\top|C \sqcap C|\exists r.C,$$

where $A$ and $r$ range over $N_C$ and $N_R$, respectively. In the following, we use symbols $A, B$ to denote atomic concepts and $C, D, E$ to denote arbitrary concepts. A *terminology* or *TBox* consists of *concept inclusion* axioms $C \sqsubseteq D$ and *concept equivalence* axioms $C \equiv D$ used as a shorthand for $C \sqsubseteq D$ and $D \sqsubseteq C$. The signature of an $\mathcal{EL}$ concept $C$ or an axiom $\alpha$, denoted by $\text{sig}(C)$ or $\text{sig}(\alpha)$, respectively, is the set of concept and role symbols occurring in it. To distinguish between the set of concept symbols and the set of role symbols, we use $\text{sig}_C(C)$ and $\text{sig}_R(C)$, respectively. The signature of a TBox $\mathcal{T}$, in symbols $\text{sig}(\mathcal{T})$ (correspondingly, $\text{sig}_C(\mathcal{T})$ and $\text{sig}_R(\mathcal{T})$), is defined analogously. Additionally, we denote the set of subconcepts occurring in a concept $C$ as $\mathtt{sub}(C)$ and the set of all subconcepts including part-conjunctions as $\mathtt{sub}_\sqcap(C)$. For instance, for $C = \exists r.(A_1 \sqcap A_2 \sqcap A_3)$ we obtain $\mathtt{sub}(C) = \{\exists r.(A_1 \sqcap A_2 \sqcap A_3), A_1 \sqcap A_2 \sqcap A_3, A_1, A_2, A_3\}$ and $\mathtt{sub}_\sqcap(C) = \{\exists r.(A_1 \sqcap A_2 \sqcap A_3), A_1 \sqcap A_2 \sqcap A_3, A_1 \sqcap A_2, A_1 \sqcap A_3, A_2 \sqcap A_3, A_1, A_2, A_3\}$. Accordingly, we denote the set of subconcepts occurring in a TBox $\mathcal{T}$ as $\mathtt{sub}(\mathcal{T})$ and the set of all subconcepts including part-conjunctions as $\mathtt{sub}_\sqcap(\mathcal{T})$.

Next, we recall the semantics of the above introduced DL constructs, which is defined by means of interpretations. An interpretation $\mathcal{I}$ is given by the domain $\Delta^\mathcal{I}$ and a function $\cdot^\mathcal{I}$ assigning each concept $A \in N_C$ a subset $A^\mathcal{I}$ of $\Delta^\mathcal{I}$ and each role $r \in N_R$ a subset $r^\mathcal{I}$ of $\Delta^\mathcal{I} \times \Delta^\mathcal{I}$. The interpretation of $\top$ is fixed to $\Delta^\mathcal{I}$. The interpretation of an arbitrary $\mathcal{EL}$ concept is defined inductively, i.e., $(C \sqcap D)^\mathcal{I} = C^\mathcal{I} \cap D^\mathcal{I}$ and $(\exists r.C)^\mathcal{I} = \{x \mid (x, y) \in r^\mathcal{I}, y \in C^\mathcal{I}\}$. An interpretation $\mathcal{I}$ satisfies an axiom $C \sqsubseteq D$ if $C^\mathcal{I} \subseteq D^\mathcal{I}$. $\mathcal{I}$ is a model of a TBox, if it satisfies all of its axioms. We say that a TBox $\mathcal{T}$ entails an axiom $\alpha$ (in symbols, $\mathcal{T} \models \alpha$), if $\alpha$ is satisfied by all models of $\mathcal{T}$. A TBox $\mathcal{T}$ entails another TBox $\mathcal{T}'$, in symbols $\mathcal{T} \models \mathcal{T}'$, if $\mathcal{T} \models \alpha$ for all $\alpha \in \mathcal{T}'$. $\mathcal{T} \equiv \mathcal{T}'$ is a shortcut for $\mathcal{T} \models \mathcal{T}'$ and $\mathcal{T}' \models \mathcal{T}$.

## 3 Reducing the Complexity of Ontologies

The size of a TBox $\mathcal{T}$ is often measured by the number of axioms contained in it ($|\mathcal{T}|$). This is, however, a simplified view of the size, which neither reflects cognitive complexity, nor the reasoning complexity. In this paper, we measure the size of a concept, an axiom, or a TBox by the number of references to signature elements as stated in the definition below.

**Definition 1.** *The* size *of an $\mathcal{EL}$ concept $D$ is defined as follows:*

- *for $D \in sig(\mathcal{T}) \cup \{\top\}$, $\int(D) = 1$;*
- *for $D = \exists r.C$, $\int(D) = \int(C) + 1$ where $r \in sig_R(\mathcal{T})$ and $C$ is an arbitrary concept;*
- *for $D = C_1 \sqcap C_2$, $\int(D) = \int(C_1) + \int(C_2)$ where $C_1, C_2$ are arbitrary concepts;*

*The* size *of an $\mathcal{EL}$ axiom (one of $C_1 \sqsubseteq C_2$, $C_1 \equiv C_2$) and a TBox $\mathcal{T}$ is accordingly defined as follows:*

- $\int(C_1 \sqsubseteq C_2) = \int(C_1) + \int(C_2)$ *for concepts $C_1, C_2$;*
- $\int(C_1 \equiv C_2) = \int(C_1) + \int(C_2)$ *for concepts $C_1, C_2$.*
- $\int(\mathcal{T}) = \sum_{\alpha \in \mathcal{T}} \int(\alpha)$ *for a TBox $\mathcal{T}$.*

The above definition, for instance, can serve as a basis for computing the average size of axioms ($\int(\mathcal{T}) \div |\mathcal{T}|$) within an ontology. In addition to the above measure of size, the number of distinct complex concept expressions $\mathtt{sub}(\mathcal{T})$ and the overall number of occurrences of such concept expressions (with the corresponding values related to $|\mathcal{T}|$) can serve as an indication of how complex are concept expressions within the ontology. In the following example, we demonstrate the difference between the two measures $|\mathcal{T}|$ and $\int(\mathcal{T})$ and show how the complexity of an ontology can be reduced in principle (by up to an exponent for ontologies without redundant axioms, i.e., axioms that can be omitted without losing any logical consequences).

*Example 2.* Let concepts $C_i$ be inductively defined by $C_0 = A$, $C_{i+1} = \exists r.C_i \sqcap \exists s.C_i$. Intuitively, $C_i$ of concepts have the shape of binary trees with exponentially many leaves. Clearly, the concepts grow exponentially with $i$, since $\int(C_i) = 2 + 2 \cdot \int(C_{i-1})$. For a natural number $n$, consider the TBox $\mathcal{T}_n$:

$$\begin{aligned} C_{n-1} &\sqsubseteq B \\ B_i &\equiv C_i \qquad 1 \leq i \leq n-1 \end{aligned}$$

While $\mathcal{T}_n$ does not contain any redundant axioms, it can easily be represented in a more compact way by recursively replacing each $C_i$ by the corresponding $B_i$, yielding $\mathcal{T}_n'$:

$$\begin{aligned} B_{n-1} &\sqsubseteq B \\ B_1 &\equiv C_1 \\ B_{i+1} &\equiv \exists r.B_i \sqcap \exists s.B_i \qquad 1 \leq i \leq n-1 \end{aligned}$$

While the number of axioms is the same in both cases, the complexity of $\mathcal{T}_n$ is clearly lower. E.g., for $n = 5$, we obtain $\int(\mathcal{T}_n) = 134$ and $\int(\mathcal{T}_n') = 24$.

We now consider the problem of finding the minimal equivalent $\mathcal{EL}$ representation for a given TBox. The corresponding decision problem can be formulated as follows:

**Definition 2 (P1).** *Given an $\mathcal{EL}$ TBox $\mathcal{T}$ and a natural number $k$, is there an $\mathcal{EL}$ TBox $\mathcal{T}'$ with $\int(\mathcal{T}') \leq k$ such that $\mathcal{T}' \equiv \mathcal{T}$.*

In general, the corresponding minimal result is not unique. We denote the set $\{\mathcal{T}' \mid \mathcal{T}' \equiv \mathcal{T}\}$ by $[\mathcal{T}]$. Note that the minimality of the result is trivially checked by deciding **P1** for a decreasing number $k$ until the answer is negative.

In literature, there are different variations of the ontology minimization problem that cover specific cases. Perhaps the simplest examples for avoidable non-succinctness are axioms that follow from other axioms and that can be removed from the ontology without losing any logical consequences. While some axioms including the last axiom in the above example can be removed in any representations, in general, subsets of axioms can be exchangeable.

*Example 3.* Consider the ontology $\mathcal{T}$:

$$C \sqsubseteq \exists r.C \qquad \exists r.D \sqsubseteq D$$
$$C \sqsubseteq D \qquad \exists r.C \sqsubseteq \exists r.D$$

$\mathcal{T}$ has two subset ontologies, $\mathcal{T}_1 = \{C \sqsubseteq \exists r.C, \exists r.C \sqsubseteq \exists r.D, \exists r.D \sqsubseteq D\}$ and $\mathcal{T}_2 = \{C \sqsubseteq \exists r.C, C \sqsubseteq D, \exists r.D \sqsubseteq D\}$. Neither of the two contains any axioms that are entailed by the remainder of the ontology. There are also no sub-expressions that can be removed. However, $\mathcal{T}_2$ is less complex than $\mathcal{T}_1$, because $C \sqsubseteq D$ is simpler (shorter) than $\exists r.C \sqsubseteq \exists r.D$.

While the above problem is already known to be non-tractable and can have many solutions, the ability to rewrite axioms of the ontology can further increase the difficulty and the number of possible solutions: While in the above cases a minimal ontology contains only subconcepts $\mathtt{sub}(\mathcal{T})$ of the original ontology $\mathcal{T}$, in general, a minimal ontology can introduce new concept expressions as demonstrated in the following example.

*Example 4.* Consider the following TBox $\mathcal{T}$:

$$C_1 \sqsubseteq A_2 \qquad A_2 \sqsubseteq C_3$$
$$\exists r.D \sqsubseteq D \qquad \exists s.C_1 \sqsubseteq D$$
$$\exists s.C_3 \sqsubseteq \exists r.(\exists s.C_1)$$

Assume that $\int(C_1)$ and $\int(C_3)$ are large. Then the axiom $\exists s.C_1 \sqsubseteq D$ needs to be exchanged by $\exists s.A_2 \sqsubseteq D$ to obtain a smaller TBox. The TBox $\mathcal{T}_m$ given below is a minimal representation of $\mathcal{T}$.

$$C_1 \sqsubseteq A_2 \qquad A_2 \sqsubseteq C_3$$
$$\exists r.D \sqsubseteq D \qquad \exists s.A_2 \sqsubseteq D$$
$$\exists s.C_3 \sqsubseteq \exists r.(\exists s.C_1)$$

We notice that the original ontology $\mathcal{T}$ does not contain the expression $\exists s.A_2 \in \mathtt{sub}(\mathcal{T}_m)$.

We can conclude that considering subsumption relations between subconcepts $\mathtt{sub}(\mathcal{T})$ of $\mathcal{T}$ is not sufficient when looking for a minimal equivalent representation. In the next section, we show that the corresponding decision problem **P1** is in fact NP-complete.

## 4 NP-Completeness

In this section, we first show the NP-hardness of the problem and then establish its NP-completeness. We show NP-hardness by a reduction from the set cover problem, which is one of the standard NP-complete problems. For a given set $\mathcal{S} = \{S_1, S_2, \ldots, S_n\}$ with carrier set $S = \bigcup_{i=1}^{n} S_i$, a *cover* $\mathcal{C} \subseteq \mathcal{S}$ is a subset of $\mathcal{S}$, such that the union of the sets in $\mathcal{C}$ covers $S$, i.e., $S = \bigcup_{C \in \mathcal{C}} C$.

The *set cover problem* is the problem to determine, for a given set $\mathcal{S} = \{S_1, S_2, \ldots, S_n\}$ and a given integer $k$, if there is a cover $\mathcal{C}$ of $\mathcal{S}$ with at most $k \geq |\mathcal{C}|$ elements. We will use a restricted version of the set cover problem, which we call the *dense set cover problem* (DSCP). In the dense set cover problem, we require that

- neither the carrier set $S$ nor the empty set is in $\mathcal{S}$,
- all singleton subsets (sets with exactly one element) of $S$ are in $\mathcal{S}$, and
- if a non-singleton set $S$ is in $\mathcal{S}$, so is some subset $S' \subseteq S$, which contains only one element less than $S$ ($|S \setminus S'| = 1$).

**Lemma 1.** *The dense set cover problem is NP-complete.*

**Proof.** Inclusion in NP is inherited from the set cover problem, of which it is a special instance.

We now reduce solving the set cover problem to solving the dense set cover problem. We start with a set cover problem for a given $\mathcal{S}$ and $k$, and first check if the carrier set $S$ is contained in $\mathcal{S}$ (if so, the problem is solved). If it is not the case, we identify the size $l$ of the largest set in $\mathcal{S}$, initialise $\mathcal{S}'$ to $\mathcal{S}$ and extend $\mathcal{S}'$ using the following algorithm:

- while $l > 1$ do
  - for all $S \in \mathcal{S}'$, choose an $s \in S$ and join $\mathcal{S}'$ with $S \setminus \{s\}$
  - decrement $l$ by one.

After this, we join $\mathcal{S}$ with $\{\{s\} \mid s \in S\}$, and remove the empty set from $\mathcal{S}$ if applicable. Note that $\mathcal{S}'$ can easily be constructed in polynomial time. Now we show that there is a cover $\mathcal{C}$ of size $\leq k$ of $\mathcal{S}$ exactly if there is a cover $\mathcal{C}'$ of size $\leq k$ of $\mathcal{S}'$. W.l.o.g., we can assume that $\emptyset \notin \mathcal{C}$, since we always obtain a cover from any cover $\mathcal{C}$ by removing $\emptyset$ from it. Since $\mathcal{S} \subseteq \mathcal{S}' \cup \{\emptyset\}$, any cover of $\mathcal{S}$ is a cover of $\mathcal{S}'$. Let $\mathcal{C}'$ be a cover of size $\leq k$ of $\mathcal{S}'$. We can construct a cover $\mathcal{C}$ of $\mathcal{S}$ by replacing each $S' \in \mathcal{C}'$ by the corresponding superset $S \in \mathcal{S}$. □

Given the above NP-completeness result, we show that the size of minimal equivalents specified in **P1** is a linear function of the size of the minimal cover. To this end, we use the lemma below to obtain a lower bound on the size of equivalents. Intuitively, it states that for each entailed non-trivial equivalence $C \equiv A$, the TBox must contain at least one axiom that is at least as large as $C' \equiv A$ for some $C'$ with $\mathcal{T} \models C \equiv C'$:

**Lemma 2.** *Let $\mathcal{T}$ be an $\mathcal{EL}$ TBox, $A \in sig(\mathcal{T})$ and $C, D$ $\mathcal{EL}$ concepts such that $\mathcal{T} \models C \equiv A$, $\mathcal{T} \models A \sqsubseteq D$ (the latter is required for induction). Then, one of the following is true:*

1. *$A$ is a conjunct of $C$ (including the case $C = A$);*
2. *there exists an $\mathcal{EL}$ concept $C'$ such that $\mathcal{T} \models C \equiv C'$ and $C' \bowtie A \in \mathcal{T}$ or $C' \bowtie A \sqcap D' \in \mathcal{T}$ for some $\bowtie \in \{\equiv, \sqsubseteq\}$ and some concept $D'$.*

**Proof.** We use the sound and complete proof system for general subsumption in $\mathcal{EL}$ terminologies introduced in [8] (Fig. 1) and prove the lemma by induction on the depth of the derivation of $C \sqsubseteq A \sqcap D$. W.l.o.g., we can assume that the proof has minimal depth. We consider the possible rules that could have been applied last to derive $C \sqsubseteq A \sqcap D$.

Basecase: $C \sqsubseteq A \sqcap D \in \mathcal{T}$ or $(\text{AX})$ was the last applied rule, in which case $C = A \sqcap D$ and condition 1 of the lemma holds.

Assume that $C \sqsubseteq A \sqcap D \notin \mathcal{T}$ and $C \neq A \sqcap D$. The rules $(\text{EX})$, and $(\text{AXTOP})$ cannot be the last rules due to the form of $C \sqsubseteq A \sqcap D$.

If $(\text{ANDL})$ was the last applied rule, then there is some sub-conjunction $D'$ of $C$ such that $\mathcal{T} \models D' \sqsubseteq A \sqcap D$. Since also $\mathcal{T} \models C \sqsubseteq D'$, we obtain $\mathcal{T} \models D' \equiv A$. We apply induction hypothesis and obtain:

1. $A$ is a conjunct of $D'$, in which case it is also a conjunct of $C$;
2. there exists an $\mathcal{EL}$ concept $C'$ such that $\mathcal{T} \models D' \equiv C'$ and $C' \bowtie A \in \mathcal{T}$ or $C' \bowtie A \sqcap D'' \in \mathcal{T}$ for some $\bowtie \in \{\equiv, \sqsubseteq\}$ and some concept $D''$. Since $\mathcal{T} \models D' \equiv C$, the second condition of the lemma holds.

If $(\text{CUT})$ was the last applied rule, then there exists a concept $E$ such that $\mathcal{T} \models C \sqsubseteq E$ and $\mathcal{T} \models E \sqsubseteq A$. Thus, $\mathcal{T} \models E \equiv A$. Once more, we apply induction hypothesis and obtain:

1. $A$ is a conjunct of $E$, in which case $E = A \sqcap E'$ for some $E'$ with $\mathcal{T} \models A \sqsubseteq E'$ and we can apply induction hypothesis to $C \sqsubseteq A \sqcap E'$, immediately obtaining the lemma;
2. there exists an $\mathcal{EL}$ concept $C'$ such that $\mathcal{T} \models E \equiv C'$ and $C' \bowtie A \in \mathcal{T}$ or $C' \bowtie A \sqcap D'' \in \mathcal{T}$ for some $\bowtie \in \{\equiv, \sqsubseteq\}$ and some concept $D''$. Since $\mathcal{T} \models E \equiv C$, the second condition of the lemma holds.

□

We now show how to encode the dense set cover problem as an ontology minimization problem. Consider an instance of the dense set cover problem with the carrier set $A = \{B_1, \ldots, B_n\}$, the set $\mathcal{S} = \{A_1, \ldots, A_m, \{B_1\}, \ldots, \{B_n\}\}$ of subsets that can be used to form a cover. By interpreting the set and element names as atomic concepts, we can construct $\mathcal{T}_{\mathcal{S}\mathsf{base}}$ as follows:

$$\mathcal{T}_{\mathcal{S}\mathsf{base}} = \{A'' \equiv A' \sqcap B \mid A'', A' \in \mathcal{S}, B \in A, A'' = A' \cup \{B\}, A'' \neq A'\}.$$

Observe that the size of $\mathcal{T}_{\mathcal{S}\mathsf{base}}$ is at least $3m$. Clearly, $\mathcal{T}_{\mathcal{S}\mathsf{base}} \models A_i \equiv \bigsqcap_{B \in A_i} B$. Let $\mathcal{T}_{\mathcal{S}} = \mathcal{T}_{\mathcal{S}\mathsf{base}} \cup \{A \equiv \bigsqcap_{B \in A} B\}$. We establish the connection between the size of $\mathcal{T}_{\mathcal{S}}$ equivalents and the size of the cover of $\mathcal{S}$ as follows:

**Lemma 3.** *$\mathcal{T}_{\mathcal{S}}$ has an equivalent of size $\int(\mathcal{T}_{\mathcal{S}\mathsf{base}}) + k + 1$ if, and only if, $\mathcal{S}$ has a cover of size $k$.*

$$\frac{}{C \sqsubseteq C}(\textsc{Ax}) \quad \frac{}{C \sqsubseteq \top}(\textsc{AxTop})$$

$$\frac{D \sqsubseteq E}{C \sqcap D \sqsubseteq E}(\textsc{AndL})$$

$$\frac{C \sqsubseteq E \quad C \sqsubseteq D}{C \sqsubseteq D \sqcap E}(\textsc{AndR})$$

$$\frac{C \sqsubseteq D}{\exists r.C \sqsubseteq \exists r.D}(\textsc{Ex})$$

$$\frac{C \sqsubseteq E \quad E \sqsubseteq D}{C \sqsubseteq D}(\textsc{Cut})$$

**Fig. 1.** Proof system for general subsumption in $\mathcal{EL}$ terminologies introduced in [8].

**Proof.** For the if-direction, assume that $\mathcal{S}$ has a cover of size $k$. We construct $\mathcal{T}'_{\mathcal{S}}$ of size $\int(\mathcal{T}_{\mathcal{S}\mathtt{base}}) + k + 1$ as follows: $\mathcal{T}'_{\mathcal{S}} = \mathcal{T}_{\mathcal{S}\mathtt{base}} \cup \{A \equiv \bigsqcap_{A' \in \mathcal{C}} A'\}$. Clearly, $\mathcal{T}'_{\mathcal{S}} \equiv \mathcal{T}_{\mathcal{S}}$.

For the only-if-direction, we assume that $k$ is minimal and argue that no equivalent $\mathcal{T}' \in [\mathcal{T}_{\mathcal{S}}]$ of size $\leq \int(\mathcal{T}_{\mathcal{S}\mathtt{base}}) + k$ can exist. Assume that $\mathcal{T}$ is a minimal TBox with $\mathcal{T} \in [\mathcal{T}_{\mathcal{S}}]$. With the observation, that the $m+n$ atomic concepts that represent elements of $\mathcal{S}$ are pairwise not equivalent with each other or the concept $A$ that represents the carrier set, we can conclude that no two atomic concepts are equivalent. From Lemma 2 it follows that, for each $A_i$ with $i \in \{1, \ldots, m\}$, there is an axiom $C_i \equiv C'_i \in \mathcal{T}$ or $C_i \sqsubseteq C'_i \in \mathcal{T}$ such that $\mathcal{T} \models C_i \equiv A_i$ and $A_i$ is a conjunct of $C'_i$ or $A_i = C'_i$. Since there are no equivalent atomic concepts and $C_i \neq A_i$ due to the minimality of $\mathcal{T}$, the size of each such axiom is at least 3 and none of these axioms coincide. Additionally, since $\mathcal{T}_{\mathcal{S}} \not\models A_i \sqsubseteq A$, $A$ cannot occur as a conjunct of $C_i$ or as a conjunct of $C'_i$;

Finally, we estimate the size of the remaining axioms and show that their cumulative size is $> k$. It also follows from Lemma 2 that there exists an axiom $C \equiv C' \in \mathcal{T}$ or $C \sqsubseteq C' \in \mathcal{T}$ such that $\mathcal{T} \models C \equiv A$ and $A$ is a conjunct of $C'$ or $A = C'$. It holds that $\mathcal{T} \models C \equiv \bigsqcap_{B \in A} B$. We also know that for no proper subset $S' \subsetneq A$ holds $\mathcal{T} \models \bigsqcap_{B \in S'} B \sqsubseteq C$. Thus, we have found a cover of $S$ and the size of the axiom must be $\geq k+1$. Thus, the overall size of $\mathcal{T}$ must be $\geq \int(\mathcal{T}_{\mathcal{S}\mathtt{base}}) + k + 1$. □

**Theorem 1. P1** *is in* NP.

**Proof.** We ask the non-deterministic algorithm to guess a TBox of the size $\leq k$. It remains to verify $\mathcal{T}' \equiv \mathcal{T}$, which can be done in PTIME [7]. □

**Theorem 2. P1** *is* NP*-complete.*

**Proof.** The problem is NP-hard as an immediate consequence of Lemmas 3 and 1. Given the result of Theorem 1, we establish NP-completeness of the problem. □

## 5 Minimizing Acyclic TBoxes

In this section, we develop an algorithm for minimizing TBoxes in polynomial time, which is guaranteed to provide a minimal TBox for a class of $\mathcal{EL}$ TBoxes satisfying a certain type of acyclicity conditions. The algorithm can also be applied to more expressive and some cyclic TBoxes, however without the guarantee of minimality.

### 5.1 Acyclicity Conditions

In this subsection, we introduce equivalence classes on concepts and discuss cyclic dependencies between equivalence classes and their impact on computing minimal representations. Let $\mathcal{T}$ be an $\mathcal{EL}$ TBox and let $C$ be a concept in $\mathtt{sub}(\mathcal{T})$. We use the notation $[C]_\mathcal{T} = \{C' \in \mathtt{sub}(\mathcal{T}) \mid \mathcal{T} \models C \equiv C'\}$ to denote the *equivalence class* of the concept $C$ and $\mathcal{C}_\mathcal{T} = \{[C]_\mathcal{T} \mid C \in \mathtt{sub}(\mathcal{T})\}$ to denote the set of all equivalence classes over the set $\mathtt{sub}(\mathcal{T})$. In case $\mathcal{T}$ is clear from the context, we omit the index. We base the acyclicity conditions on the following reference relations, which use both syntactic and semantic dependencies between equivalence classes:

**Definition 3.** *Let $\mathcal{T}$ be an $\mathcal{EL}$ TBox. The* reference relations $\prec_\sqsubseteq$, $\prec_\sqsupseteq$ *and* $\prec_s$*, all subsets of $\mathcal{C} \times \mathcal{C}$, are given as follows:*

- $[C] \prec_s [C']$ *if, for some $C_1 \in [C], C_2 \in [C']$, it holds that $C_2$ occurs in $C_1$;*
- $[C] \prec_\sqsubseteq [C']$ *if, for some $C_1 \in [C], C_2 \in [C']$, it holds that $[C_1] \prec_s [C_2]$ or $\mathcal{T} \models C_1 \sqsubseteq C_2$;*
- $[C] \prec_\sqsupseteq [C']$ *if, for some $C_1 \in [C], C_2 \in [C']$, it holds that $[C_1] \prec_s [C_2]$ or $\mathcal{T} \models C_1 \sqsupseteq C_2$.*

We call a TBox *cyclic*, if any of the above relations $\prec_\sqsubseteq, \prec_\sqsupseteq, \prec_s$ is cyclic. We say that a TBox $\mathcal{T}$ is *strongly cyclic* if $\prec_s$ is cyclic. The algorithm presented in this paper is applicable for TBoxes not containing strong cycles. Most of the large bio-medical ontologies including Galen, Gene Ontology and NCI do not contain strong cycles. This was also the case for earlier versions of SNOMED, e.g., the one dated 09 February 2005 [9]. Note that asking for the absence of cycles in $\prec_s$ is a weaker requirement than for $\prec_\sqsubseteq$ or $\prec_\sqsupseteq$, as $\prec_s \subseteq \prec_\sqsubseteq \cap \prec_\sqsupseteq$. But the reverse relationship between the conditions holds. In the remainder of this section, we use the following sufficient but not necessary criterion for cyclicity.

**Lemma 4.** *Let $\mathcal{T}$ be an $\mathcal{EL}$ TBox and $C, C'$ two concepts such that $C'$ contains $C$ within the scope of an existential restriction and there is $D \in \mathtt{sub}(\mathcal{T})$ with $\mathcal{T} \models D \equiv C$. Then $\mathcal{T}$ is cyclic, if $\mathcal{T} \models C \sqsubseteq C'$ or $\mathcal{T} \models C' \sqsubseteq C$ and no conjuncts can be removed from $C'$ without invalidating the statement $\mathcal{T} \models C' \sqsubseteq C$.*

**Proof Sketch.** If $C \notin \mathtt{sub}(\mathcal{T})$, we obtain the witness reference path for cyclicity by considering the subsumption between $D$ and the concept $D'$, which is obtained from $C'$ by replacing $C$ by $D$. We consider the grammars introduced in [8]. Since $D \in \mathtt{sub}(\mathcal{T})$, there is a corresponding initial non-terminal, starting from which we can generate $D'$. Since each subsumee (subsumer) grammar transition corresponds to a subsumee (subsumer) reference relation, the derivation of $D'$ in the grammar is a witness reference path for cyclicity. □

In some cases, TBoxes contain cycles that are caused by redundant conjuncts and can easily be removed.

*Example 5.* $\{A \sqcap B \sqsubseteq C, A \sqsubseteq B\}$ has a cyclic $\prec_{\sqsupseteq}$ relation due to a cycle between $A \sqcap B$ and $A$. It can be transformed into an acyclic TBox $\{A \sqsubseteq C, A \sqsubseteq B\}$.

We call conjunctions $C' \sqcap C''$ in $\mathtt{sub}(\mathcal{T})$ such that $\mathcal{T} \models C' \sqsubseteq C' \sqcap C''$ *subsumer-containing conjunctions*. We can easily eliminate subsumer-containing conjunctions in TBoxes before applying the algorithm: for each subsumer-containing conjunction $C' \sqcap C''$ in $\mathtt{sub}(\mathcal{T})$ with $\mathcal{T} \models C' \sqsubseteq C' \sqcap C''$, we replace $C' \sqcap C''$ in $\mathcal{T}$ by $C'$, and add the axiom $C' \sqsubseteq C''$ to $\mathcal{T}$. We can show that the closure of each equivalence class $[C]$ of an acyclic TBox $\mathcal{T}$ is finite if we exclude subsumer-containing conjunctions. We denote such a closure with $[C]^* = \{C' \mid \mathcal{T} \models C \equiv C'$ and $C'$ is not a subsumer-containing conjunction$\}$. It can be shown that, for an acyclic TBox, each $[C]^*$ is free of conjunctions. We denote the extended set of subconcepts of $\mathcal{T}$ by $\mathtt{sub}(\mathcal{T})^* = \bigcup_{[C] \in \mathcal{C}} [C]^*$.

Another kind of removable cyclic dependencies are conjunctions on the right-hand side. We use a simple *decomposition*, in which all conjunctions on the right-hand side of axioms are replaced by separate inclusion axioms for each conjunct. We obtain the decomposed version $\mathcal{T}'$ of a TBox $\mathcal{T}$ by replacing each $C \sqsubseteq D_1 \sqcap D_2 \in \mathcal{T}_m$ by $C \sqsubseteq D_1, C \sqsubseteq D_2$ until a fixpoint is reached. *Composition* is the dual transformation: we replace any two axioms $C \sqsubseteq D_1, C \sqsubseteq D_2$ by $C \sqsubseteq D_1 \sqcap D_2$ until a fixpoint is reached.

Unless we state otherwise, we assume in the following that TBoxes are decomposed and do not contain subsumer-containing conjunctions.

### 5.2 Uniqueness of Minimal TBoxes

Acyclic TBoxes are better behaved not only with respect to the complexity of minimization, but they also have a unique minimal TBox modulo replacement of equivalent concepts by one another (if we assume that the TBox with the lower number of equivalence axioms should be preferred in case of equally large TBoxes).

To be able to determine a unique syntactic representation of a TBox $\mathcal{T}$, we choose a representative $C' \in [C]^*$ for each equivalence class $[C] \in \mathcal{C}$ and denote it using the *representative selection function* $r : \mathcal{C} \to \mathtt{sub}(\mathcal{T})^*$ with $r([C]) = C'$. We say that $r$ is *valid*, if for all $[C], [D] \in \mathcal{C}$ with $[C] \neq [D]$ it holds that $C' \in [C]^*$ occurs in $r([D])$ only if $C' = r([C])$, i.e., representatives can only contain other representatives, but not other elements of equivalence classes.

Let $\bowtie \in \{\equiv, \sqsubseteq\}$. We say that $\mathcal{T}$ is *aligned* with $r$, if for each $C \bowtie D \in \mathcal{T}$ one of the following conditions hold:

- if $\mathcal{T} \not\models C \equiv D$, then $C = r([C])$ and $D = r([D])$;
- if $\mathcal{T} \models C \equiv D$, then for each $C'$ such that $C' \neq C$, $C' \neq D$ and $C'$ occurs in $C$ or $D$ it holds that $C' = r([C'])$.

In other words, the only axioms, in which we allow an occurrence of a non-representative $C$ are axioms relating $C$ with concepts equivalent to it.

Since minimal TBoxes can sometimes contain subsumption axioms relating two equivalent concepts with each other, the otherwise unique TBox result can vary in the choice between subsumption and equivalence axioms. For the sake of uniqueness, we assume that, whenever we have a choice between equivalence ($\equiv$) and subsumption axioms ($\sqsubseteq$) in the resulting TBox, we prefer subsumption axioms.

We call a TBox *non-redundant*, if there is no $\alpha \in \mathcal{T}$ such that $\mathcal{T} \setminus \{\alpha\} \models \alpha$. In order to show how to compute a minimal equivalent TBox for an acyclic initial TBox, we first show that we do not need new equivalence classes or new relations between them to obtain any non-redundant, decomposed, equivalent TBox. In other words, non-redundant, decomposed axioms encoding relations between equivalence classes are unique up to exchanging equivalent concepts.

**Lemma 5.** *Let $\mathcal{T}_1, \mathcal{T}_2$ be two non-redundant, acyclic $\mathcal{EL}$ TBoxes such that $\mathcal{T}_1 \equiv \mathcal{T}_2$. Let $C \sqsubseteq D \in \mathcal{T}_2$ such that $\mathcal{T}_2 \not\models C \equiv D$. Then there is $C' \sqsubseteq D' \in \mathcal{T}_1$ such that $\mathcal{T}_1 \models C' \equiv C$, $\mathcal{T}_1 \models D' \equiv D$.*

**Proof.** Since $\mathcal{T}_1 \equiv \mathcal{T}_2$, it holds that $\mathcal{T}_1 \models C \sqsubseteq D$. We prove the lemma by induction on the structure of a proof of $C \sqsubseteq D$ from $\mathcal{T}_1$ using the $\mathcal{EL}$ deduction calculus in Fig.1. W.l.o.g., we assume that proofs have minimal depth. We can show that, within this proof, there is a sequent $C' \sqsubseteq D'$ with the following properties:

a. $\mathcal{T}_1 \models C \equiv C'$ and $\mathcal{T}_1 \models D \equiv D'$;
b. either $C' \sqsubseteq D'$ is a leaf within the derivation or the sequents used to derive $C' \sqsubseteq D'$ do not satisfy [a];
c. the last rule applied to derive $C' \sqsubseteq D'$ was not EX;
d. the last rule applied to derive $C' \sqsubseteq D'$ was not AX;
e. $C' \sqsubseteq D'$ can only be derived from $\mathcal{T}_2$ using the axiom $C \sqsubseteq D$.

Such a sequent exists for the following reasons:

The sequent $C \sqsubseteq D$ is part of the proof. Thus, there is always at least one sequent satisfying [a]. Since the derivation is a finite tree, we can show that there is always a sequent satisfying [a] and [b].

For [c], [d], [e], let us assume that we have a sequent $C' \sqsubseteq D'$ for which [a] and [b] hold, but not [c] or not [d] or not [e]. In order to show the opposite, we consider the rest of the derivation, namely the derivation path between the two sequents $C \sqsubseteq D$ and $C' \sqsubseteq D'$. In principle, this derivation path can involve ANDL, CUT, EX and ANDR. Since both TBoxes are acyclic, they cannot contain any conjunctions. Thus, we can assume that our depth-minimal proof does not use ANDL or ANDR. If EX is involved somewhere on this derivation path, then the TBox is cyclic by Lemma 4 due to $\mathcal{T}_1 \models C \equiv C'$ and $\mathcal{T}_1 \models D \equiv D'$; (we can construct a witness concepts containing $C'$ and $D'$ within existential restrictions by removing each CUT from the derivation path, i.e., replacing concepts on the right (left) by the concepts on the left (right) throughout the derivation path until it consists of EX only; we then find the witness concepts in the place of $C$ and $D$). Thus, we can exclude EX from this derivation path. As a consequence, the derivation path only contains CUT. We consider sequents used in CUT rules: there is a sequence of concepts $C_1, \ldots, C_n$ and $D_1, \ldots, D_m$ with $\mathcal{T}_2 \models C_i \sqsubseteq C_{i+1}$, $\mathcal{T}_2 \models C \sqsubseteq C_1$, $\mathcal{T}_2 \models C_n \sqsubseteq C'$ and $\mathcal{T}_2 \models D_{i+1} \sqsubseteq D_i$, $\mathcal{T}_2 \models D' \sqsubseteq D_1$, $\mathcal{T}_2 \models D_m \sqsubseteq D$. Since $\mathcal{T}_1 \equiv \mathcal{T}_2$, each of the above subsumptions is also derivable in $\mathcal{T}_2$. If all of the subsumptions above can be derived from $\mathcal{T}_2$ without using the axiom $C \sqsubseteq D$, then $\mathcal{T}_2$ is redundant because we can remove $C \sqsubseteq D$ from $\mathcal{T}_2$

without losing any consequences. Thus, at least one of the above subsumptions must only be derivable from $\mathcal{T}_2$ using $C \sqsubseteq D$, i.e., it must satisfy [e].

We now show that if one of [c], [d] does not hold for one of the avobe subsumptions, then nor does [e]. First, we consider $C' \sqsubseteq D'$. In case of AX, we have $C' = D'$, which can be derived in any TBox and, thus, without using $C \sqsubseteq D$. In case of EX, $C' = \exists r.C''$ and $D' = \exists r.D''$ for some role $r$ and concepts $C'', D''$ with $\mathcal{T}_1 \models C'' \sqsubseteq D''$. Assume for a contradiction that $C'' \sqsubseteq D''$ cannot be derived in $\mathcal{T}_2$ without using $C \sqsubseteq D$. Also here, we can consider the depth-minimal derivation path from $C \sqsubseteq D$ to $C'' \sqsubseteq D''$. The path can be extended by applying EX to derive $C' \sqsubseteq D'$. Again, we can construct witness concepts for cyclicity of $\mathcal{T}$ by Lemma 4 containing $C$ and $D$ within the scope of an existential restriction (by excluding ANDL or ANDR and removing CUT from the derivation path). We conclude that [e] can only hold if all three conditions ([c],[d],[e]) hold.

Now, we show the claim for the subsumptions on the left or right of $C' \sqsubseteq D'$. Assume that one of the sequents left from $C'$ or right from $D'$ is only derivable from $\mathcal{T}_2$ using $C \sqsubseteq D$, e.g., $C_i \sqsubseteq C_{i+1}$. Once again we consider the corresponding derivation path ($\mathcal{T}_2$). Since both TBoxes are acyclic, they cannot contain any conjunctions. Thus, we can assume that our depth-minimal proof does not use ANDL or ANDR. We can exclude EX for the same reasons as above: If EX is involved somewhere on this derivation path, then we can construct corresponding subsumee and subsumer concepts containing $C$ and $D$ within the scope of existential restrictions. Then, the TBox $\mathcal{T}_2$ would be cyclic by Lemma 4. We can conclude that $\mathcal{T}_2 \models D \sqsubseteq C_{i+1}$ and $\mathcal{T}_2 \models C_i \sqsubseteq C$. Thus, we have found a sequent, which satisfies [a], [b] and [d],[e]. If the last applied rule for deriving $C_i \sqsubseteq C_{i+1}$ was EX, then both TBoxes are cyclic: by assumption, $C_i \sqsubseteq C_{i+1}$ is only derivable from $\mathcal{T}_2$ using $C \sqsubseteq D$. Then, also the preceding sequent (before EX) can only be derived from $\mathcal{T}_2$ using the axiom $C \sqsubseteq D$. If this is the case, we can again construct witness concepts containing $C$ and $D$ within the scope of existential restrictions. Thus, $C_i \sqsubseteq C_{i+1}$ also satisfies [c]. If we instead assume that one of the sequents right from $D'$ can only be derived from $\mathcal{T}_2$ using the axiom $C \sqsubseteq D$, then the argumentation is the same as above to show that it satisfies all five conditions.

Let $C' \sqsubseteq D'$ be such a sequent that satisfies [a], [b], [c],[d] and [e]. We now show that $C' \sqsubseteq D' \in \mathcal{T}_1$ by excluding the remaining cases. We consider the last applied rules of the derivation of $C' \sqsubseteq D'$. We have excluded EX, AX. We can exclude ANDL or ANDR, since both TBoxes are acyclic and they cannot contain any conjunctions. We can exclude CUT: assume $C' \sqsubseteq C''$ and $C'' \sqsubseteq D'$ are the sequents before $C' \sqsubseteq D'$. If we assume that $C' \sqsubseteq D'$ is only derivable using $C \sqsubseteq D$ and $\mathcal{T}_1$ is not cyclic, we can show that either $\mathcal{T}_1 \models C'' \equiv D'$ or $\mathcal{T}_1 \models C'' \equiv C'$. Thus, [b] does not hold as assumed. We can also exclude AXTOP due to [e] and the non-redundancy of $\mathcal{T}_2$. □

While the above lemma addresses relations between equivalence classes in non-redundant, decomposed TBoxes, it does not allow us to draw conclusions about axioms representing relations within equivalence classes. The purpose of the below lemma is to determine the part of the TBox that encodes relations between equivalent concepts within equivalence classes. For this, we divide the TBox $\mathcal{T}$ into partitions: one for non-equivalence axioms $\mathcal{T}^0 = \{C \sqsubseteq D \in \mathcal{T} \mid \mathcal{T} \not\models C \equiv D\}$ and one for axioms encoding relations within each equivalence class: $\mathcal{T}^{[C']} = \{C \equiv D \in \mathcal{T} \mid C, D \in [C']\}$ for each

$[C'] \in \mathcal{C}$. We denote the set of all subsumption dependencies holding within a partition by $\mathcal{T}^{\texttt{full},[C']} = \{C \sqsubseteq D \mid C, D \in [C']\}$. In each (equivalence class) partition, a part of dependencies can be deducible from the remainder of the TBox.

*Example 6.* Consider the TBox $\mathcal{T} = \{A \sqsubseteq B, \exists r.A \equiv \exists r.B\}$. For the equivalence class $\{\exists r.A, \exists r.B\}$, the subsumption $\exists r.A \sqsubseteq \exists r.B$ follows from $A \sqsubseteq B$.

We denote entailed dependencies for an equivalence class $[C']$ by $\mathcal{T}^{\texttt{red},[C']} = \{C \sqsubseteq D \in \mathcal{T}^{\texttt{full},[C']} \mid \mathcal{T} \setminus \mathcal{T}^{\texttt{full},[C']} \models C \sqsubseteq D\}$. We now consider alternative representations of each partition $\mathcal{T}^{[C']}$, all of which are subsets of $\mathcal{T}^{\texttt{full},[C']}$. We first show that, in any acyclic TBox $\mathcal{T}$ aligned with some valid $r$, we can determine the entailed dependencies $\mathcal{T}^{\texttt{red},[C']}$ within each $\mathcal{T}^{\texttt{full},[C']}$ based on $\mathcal{T}^0$.

**Lemma 6.** *Let $\mathcal{T}$ be a non-redundant, acyclic $\mathcal{EL}$ TBox aligned with a valid representative selection function $r$. Then, for each non-singleton equivalence class $[C''] \in \mathcal{C}(\mathcal{T})$ and each pair $C, D \in [C'']$, it holds that $C \sqsubseteq D \in \mathcal{T}^{\texttt{red},[C'']}$ exactly if one of the following conditions is true:*

1. *$D = \top$*
2. *there are concepts $C', D'$ such that $C = \exists r.C', D = \exists r.D'$ and $\mathcal{T} \models C' \sqsubseteq D'$, $\mathcal{T} \not\models C' \equiv D'$.*

**Proof.** We start with the if-direction. Assume that $D = \top$. For each $C$ and each $\mathcal{T}$, it always holds that $\mathcal{T} \setminus \mathcal{T}^{\texttt{full},[C'']} \models C \sqsubseteq D$. Thus, also $C \sqsubseteq D \in \mathcal{T}^{\texttt{red},[C'']}$. Assume that the second condition holds: there are concepts $C', D'$ such that $C = \exists r.C', D = \exists r.D'$ and $\mathcal{T} \models C' \sqsubseteq D'$, $\mathcal{T} \not\models C' \equiv D'$. Since $\mathcal{T} \not\models C' \equiv D'$, there is an axiom $r[C'] \sqsubseteq r[D'] \in \mathcal{T}^0$. Additionally $\mathcal{T}^{[C']} \models r[C'] \equiv C'$ and $\mathcal{T}^{[D']} \models r[D'] \equiv D'$. Thus, $\mathcal{T}^{[C']} \cup \mathcal{T}^{[D']} \cup \mathcal{T}^0 \models C \sqsubseteq D$ and, due to acyclicity of $\mathcal{T}$, $(\mathcal{T}^{[C']} \cup \mathcal{T}^{[D']} \cup \mathcal{T}^0) \setminus \mathcal{T}^{\texttt{full},[C'']} \models C \sqsubseteq D$. Therefore, $C \sqsubseteq D \in \mathcal{T}^{\texttt{red},[C'']}$.

For the only-if-direction, we assume that $C \sqsubseteq D \in \mathcal{T}^{\texttt{red},[C'']}$ and show that one of the above two cases is true. To be able to use induction, we prove a more general version of the statement: we show that one of the above two cases is true for all $C \sqsubseteq D \in \mathcal{T}^{\texttt{red},[C'']^*} = \{C \sqsubseteq D \in \mathcal{T}^{\texttt{full},[C']^*} \mid \mathcal{T} \setminus \mathcal{T}^{\texttt{full},[C']^*} \models C \sqsubseteq D\}$ ($C \sqsubseteq D \in \mathcal{T}^{\texttt{red},[C'']}$ is a special case thereof). Let $C \sqsubseteq D \in \mathcal{T}^{\texttt{red},[C'']^*}$. Then, $\mathcal{T} \setminus \mathcal{T}^{\texttt{full},[C']^*} \models C \sqsubseteq D$. For brevity, we denote $\mathcal{T} \setminus \mathcal{T}^{\texttt{full},[C']^*}$ by $\mathcal{T}_1$. We prove the claim by induction on the structure of a proof of $C \sqsubseteq D$ from $\mathcal{T}_1$ using the $\mathcal{EL}$ deduction calculus in Fig.1. W.l.o.g., we assume that proofs have minimal depth. We start with three cases that are the basis of inducting.

- We can exclude the case $C \sqsubseteq D \in \mathcal{T}_1$, since $C \sqsubseteq D \in \mathcal{T}^{\texttt{full},[C']^*}$.
- Since $C, D \in [C'']^*$ are assumed to be different, we can also exclude AX.
- If the sequent has been derived by applying AXTOP, then $D = \top$, which corresponds to the first case of our lemma.

Since $\mathcal{T}$ is acyclic, $C, D$ and the TBox itself cannot contain any conjunctions. Thus, we can assume that our depth-minimal proof does not use ANDL or ANDR. It remains to consider CUT and EX.

CUT In this case, there is a concept $C'$ such that $\mathcal{T} \models C \sqsubseteq C'$ and $\mathcal{T} \models C' \sqsubseteq D$. We conclude that $\mathcal{T} \models C \equiv C'$ and $\mathcal{T} \models C' \equiv D$. Both of these sequents are also in $\mathcal{T}^{\mathtt{red},[C'']^*}$, since they follow from $\mathcal{T}_1$. By induction hypothesis, we can conclude that one of the two above cases holds for each of them. However, the only valid combination of cases is when the second condition holds for both sequents. Thus, there are concepts $C_1', D_1', D_2'$ such that $C = \exists r.C_1', C' = \exists r.D_1', D = \exists r.D_2'$ and $\mathcal{T} \models C_1' \sqsubseteq D_1', \mathcal{T} \models D_1' \sqsubseteq D_2'$. Additionally, it holds that $\mathcal{T} \not\models C_1' \equiv D_1'$ and $\mathcal{T} \not\models D_1' \equiv D_2'$. It follows that $\mathcal{T} \models C_1' \sqsubseteq D_2'$ and $\mathcal{T} \not\models C_1' \equiv D_2'$. Thus, the second case is true for $C, D$.

EX In this case, there are concepts $C', D'$ such that $C = \exists r.C', D = \exists r.D'$ and $\mathcal{T} \models C' \sqsubseteq D'$. It remains to show that $\mathcal{T} \not\models C' \equiv D'$. Since $\mathcal{T}$ is aligned with $r$ and $\mathcal{T} \models C' \equiv D'$, it follows that $C, D$ cannot be both in $[C'']$ at the same time unless $r[C'] = C' = D'$. The latter case is not possible due to $C \neq D$.

□

As a consequence, each equivalence class partition can be considered independently from other equivalence class partitions. In particular, this implies that, for any syntactic representation $\mathcal{T}^{[C]}$ of a partition for equivalence class $[C]$, we can obtain $\mathcal{T}^{\mathtt{full},[C]}$ from $\mathcal{T}^{[C]} \cup \mathcal{T}^{\mathtt{red},[C]}$ by computing its transitive closure [4].

**Lemma 7.** *Let $\mathcal{T}$ be a non-redundant, acyclic $\mathcal{EL}$ TBox aligned with a valid representative selection function $r$. Then, for each equivalence class $[C] \in \mathcal{C}(\mathcal{T})$ it holds that $(\mathcal{T}^{[C]} \cup \mathcal{T}^{\mathtt{red},[C]})^* = \mathcal{T}^{\mathtt{full},[C]}$.*

**Proof.** $(\mathcal{T}^{[C]} \cup \mathcal{T}^{\mathtt{red},[C]})^* \subseteq \mathcal{T}^{\mathtt{full},[C]}$ follows from the transitivity of $\sqsubseteq$ and the inclusion $\mathcal{T}^{[C]} \cup \mathcal{T}^{\mathtt{red},[C]} \subseteq \mathcal{T}^{\mathtt{full},[C]}$. For $(\mathcal{T}^{[C]} \cup \mathcal{T}^{\mathtt{red},[C]})^* \supseteq \mathcal{T}^{\mathtt{full},[C]}$, we consider an inclusion $C \sqsubseteq D \in \mathcal{T}^{\mathtt{full},[C]}$ and show by induction on the structure of a depth-minimal derivation $\mathcal{T} \vdash C \sqsubseteq D$ that $C \sqsubseteq D \in (\mathcal{T}^{[C]} \cup \mathcal{T}^{\mathtt{red},[C]})^*$. We start with three cases that are the basis of inducting.

- In the case that $C \sqsubseteq D \in \mathcal{T}$, $C \sqsubseteq D \in \mathcal{T}^{[C]}$ and therefore in $(\mathcal{T}^{[C]} \cup \mathcal{T}^{\mathtt{red},[C]})^*$.
- Since $C, D \in [C]$ are assumed to be different, we can exclude AX.
- If the sequent has been derived by applying AXTOP, then $D = \top$. It follows that $C \sqsubseteq D \in \mathcal{T}^{\mathtt{red},[C]}$ by Lemma 6 and therefore $C \sqsubseteq D$ is in $(\mathcal{T}^{[C]} \cup \mathcal{T}^{\mathtt{red},[C]})^*$.

Since $\mathcal{T}$ is acyclic, $C, D$ and the TBox itself cannot contain any conjunctions. Thus, we can assume that our depth-minimal proof does not use ANDL or ANDR. It remains to consider CUT and EX.

EX: In this case, there are concepts $C', D'$ such that $C = \exists r.C', D = \exists r.D'$ and $\mathcal{T} \models C' \sqsubseteq D'$. It remains to show that $\mathcal{T} \not\models C' \equiv D'$. Since $\mathcal{T}$ is aligned with $r$ and $\mathcal{T} \models C' \equiv D'$, it follows that $C, D$ cannot be both in $[C'']$ at the same time unless $r[C'] = C' = D'$. The latter case is not possible due to $C \neq D$. It follows that $C \sqsubseteq D \in \mathcal{T}^{\mathtt{red},[C]}$ by Lemma 6 and therefore $C \sqsubseteq D$ is in $(\mathcal{T}^{[C]} \cup \mathcal{T}^{\mathtt{red},[C]})^*$.

[4] For a set $\mathcal{T}$ of axioms, the transitive closure $(\mathcal{T})^*$ is obtained by including $C \sqsubseteq D$ for any $C, D$ such that there exists $C'$ with $\mathcal{T} \models \{C \sqsubseteq C', C' \sqsubseteq D\}$.

CUT: In this case, we apply induction hypothesis: there is a concept $C'$ such that $\mathcal{T} \models C \sqsubseteq C'$ and $\mathcal{T} \models C' \sqsubseteq D$. We conclude that $\mathcal{T} \models C \equiv C'$ and $\mathcal{T} \models C' \equiv D$. Additionally, it can be shown that, if $C, D \in \texttt{sub}(\mathcal{T})$, any sequent involved in a depth-minimal proof of $C \sqsubseteq D$ from $\mathcal{T}$ is also in $\texttt{sub}(\mathcal{T})$. Thus, both of these sequents are in $(\mathcal{T}^{[C]} \cup \mathcal{T}^{\texttt{red},[C]})^*$ by induction hypothesis. Due to transitivity of $\sqsubseteq$, we conclude that also $C \sqsubseteq D$ is in $(\mathcal{T}^{[C]} \cup \mathcal{T}^{\texttt{red},[C]})^*$.

□

Since our implementation operates on ontologies represented in the OWL Web Ontology Language, we consider here an important detail of this language. In addition to constructs mentioned in preliminaries, OWL Web Ontology Language allows for *OwlEquivalentClassesAxioms* - axioms, in which we can specify a set of equivalent concepts. With the exception of equivalence classes containing $\top$, for which there exists an equally small representation without an OwlEquivalentClassesAxiom, this is clearly the smallest representation for equivalence class partitions.

Let $[C]^{\texttt{nonred}} = [C] \setminus \{C' \in [C] \mid C' \sqsubseteq D_1' \text{ and } C' \sqsupseteq D_2' \in \mathcal{T}^{\texttt{red},[C]}$ for some $D_1', D_2'\}$. Let $\mathcal{T}^{\texttt{nonred},[C]}$ be the corresponding OWLEquivalentClassesAxiom with $[C]^{\texttt{nonred}}$ as the set of equivalent concepts. Note that, according to the semantics of OwlEquivalentClassesAxioms, it holds that $\mathcal{T}^{\texttt{nonred},[C]} \models \mathcal{T}^{\texttt{full},[C]^{\texttt{nonred}}}$. Thus, $\mathcal{T}^{\texttt{nonred},[C]} \cup \mathcal{T}^{\texttt{red},[C]} \models \mathcal{T}^{\texttt{full},[C]}$. Note that $\int(\mathcal{T}^{\texttt{nonred},[C]}) = \sum_{C' \in [C]^{\texttt{nonred}}} \int(C')$.

**Lemma 8.** *Let $\mathcal{T}$ be a non-redundant, acyclic $\mathcal{EL}$ TBox aligned with a valid representative selection function $r$. Then, $\int(\mathcal{T}^{\texttt{nonred},[C]}) \leq \int(\mathcal{T}^{[C]})$ for each equivalence class $[C] \in \mathcal{C}(\mathcal{T})$.*

**Proof Sketch.** $\int(\mathcal{T}^{\texttt{nonred},[C]}) = \sum_{C' \in [C]^{\texttt{nonred}}} \int(C')$. As for $\int(\mathcal{T}^{[C]})$, we can show using Lemma 7 that each $C' \in [C]^{\texttt{nonred}}$ occurs in it at least once. Additionally, we can show that there are no $C_1, C_2 \in [C]$ such that $C_1$ occurs in $C_2$ due to the acyclicity of $\mathcal{T}$. Thus, $\int(\mathcal{T}^{[C]}) \geq \sum_{C' \in [C]^{\texttt{nonred}}} \int(C')$. □

Based on the above lemmas, we can show that, in the acyclic case, we can compute a minimal TBox by eliminating redundant axioms, fixing the representative selection function $r$ to some minimal value, constructing the core representation $\mathcal{T}^{\texttt{nonred},[C]}$ for each non-singleton equivalence class $[C]$ and composing $\mathcal{T}$ again. We say that $r$ is *minimal*, if for each $[C] \in \mathcal{C}$ holds: there is no $C' \in [C]^*$ such that $\int(C') < \int(r([C]))$.

**Theorem 3.** *Let $\mathcal{T}$ be a non-redundant, acyclic $\mathcal{EL}$ TBox and $r$ a minimal, valid representative selection function. Let the TBox $\mathcal{T}_n = \mathcal{T}^0 \cup \bigcup_{[C] \in \mathcal{C}, |[C]| \geq 2} \mathcal{T}^{\texttt{nonred},[C]}$ be aligned with $r$. Let $\mathcal{T}_n'$ be a composed version of $\mathcal{T}_n$. Then, for any minimal TBox $\mathcal{T}_m$ with $\mathcal{T}_m \equiv \mathcal{T}$ it holds that $\int(\mathcal{T}_m) = \int(\mathcal{T}_n')$.*

**Proof Sketch.** We decompose $\mathcal{T}_m$ and obtain $\mathcal{T}_d$. We now divide $\mathcal{T}_d$ into partitions $\mathcal{T}_d^0$,$\mathcal{T}_d^{[C_i]}$ for $C_i \in \mathcal{C}(\mathcal{T}_d)$. The following hold:

- $\mathcal{T}_d$ contains no conjunctions.
- For each $C \sqsubseteq D \in \mathcal{T}_d^0$, there is $C' \sqsubseteq D' \in \mathcal{T}_n^0$ such that $\mathcal{T} \models C \equiv C'$ and $\mathcal{T} \models D \equiv D'$ and it holds that $\int(C \sqsubseteq D) \geq \int(C' \sqsubseteq D')$.
- Every acyclic non-redundant TBox $\mathcal{T}'$ has exactly one minimal composed version, i.e., a representation obtained by composing axioms with equal terms on the left-hand side until no compositions are possible.

**Algorithm 1:** Rewriting $\mathcal{T}_{\texttt{in}}$

**Data**: $\mathcal{T}_{\texttt{in}}$ acyclic decomposed TBox
**Result**: $\mathcal{T}_{\texttt{out}}$ minimal equivalent TBox

1 $\mathcal{C}_{\texttt{all}} \leftarrow \mathcal{C}$;
2 $\mathcal{C}_{\texttt{TODO}} \leftarrow \mathcal{C}_{\texttt{all}}$;
3 $\mathcal{T}_{\texttt{out}} \leftarrow$ remove equivalence axioms from $\mathcal{T}_{\texttt{in}}$;
4 **while** $\mathcal{C}_{\texttt{TODO}} \neq \emptyset$ **do**
5     **for** $[C] \in$ *leaves*$(\mathcal{C}_{\texttt{TODO}})$ **do**
6         choose minimal representative $r([C])$;
7         replace $C' \in [C]$ in $\mathcal{T}_{\texttt{out}}$ by $r([C])$;
8         replace $C' \in [C]$ in $\mathcal{C}_{\texttt{TODO}} \setminus \{[C]\}$ by $r([C])$;
9         replace $C' \in [C]$ in $\mathcal{C}_{\texttt{all}} \setminus \{[C]\}$ by $r([C])$;
10         $\mathcal{C}_{\texttt{TODO}} \leftarrow \mathcal{C}_{\texttt{TODO}} \setminus \{[C]\}$;
11 $\mathcal{T}_e \leftarrow \bigcup_{[C] \in \mathcal{C}_{\texttt{all}}, |[C]| \geq 2} \mathcal{T}^{\texttt{nonred},[C]}$;
12 **for** $\alpha \in \mathcal{T}_{\texttt{out}}$ **do**
13     **if** $\mathcal{T}_{\texttt{out}} \cup \mathcal{T}_e \setminus \{\alpha\} \models \alpha$ **then**
14         $\mathcal{T}_{\texttt{out}} \leftarrow \mathcal{T}_{\texttt{out}} \setminus \{\alpha\}$;
15 $\mathcal{T}_{\texttt{out}} \leftarrow \mathcal{T}_{\texttt{out}} \cup \mathcal{T}_e$ ;
16 $\mathcal{T}_{\texttt{out}} \leftarrow$ compose$(\mathcal{T}_{\texttt{out}})$;

– Composition affects only axioms from $\mathcal{T}_d^0$ and $\mathcal{T}_n^0$.
– After composition, it holds that $|\mathcal{T}_m^0| \geq |\mathcal{T}_n'^0|$ (due to alignment with $r$).
– After composition, it holds that $\int(\mathcal{T}_m^0) \geq \int(\mathcal{T}_n'^0)$.
– For each $\mathcal{T}_d^{[C]}$ with $C \in \mathcal{C}(\mathcal{T}_d)$ holds: $\int(\mathcal{T}_d^{[C]}) \geq \int(\mathcal{T}^{\texttt{nonred},[C]})$.
– $\int(\mathcal{T}_m) \geq \int(\mathcal{T}_n')$.

Due to our assumption of minimality for $\mathcal{T}_m$, it follows that $\int(\mathcal{T}_m) = \int(\mathcal{T}_n')$. □

Algorithm 1 implements the iterative computation of (a minimal) $r$ and the minimal TBox $\mathcal{T}_n'$. It takes an acyclic decomposed TBox $\mathcal{T}_{\texttt{in}}$ and computes the corresponding minimal equivalent TBox $\mathcal{T}_{\texttt{out}}$. Line 3 is not strictly necessary, but allows for a more efficient processing. In Lines 4-10, a minimal representative selection function $r$ is iteratively determined – for one equivalence class at a time – and all data structures are aligned with $r$. We distinguish two versions of equivalence classes: $\mathcal{C}_{\texttt{TODO}}$ contains equivalence classes, for which the minimal representative has not been selected yet. In each iteration, we process the leaves in $\mathcal{C}_{\texttt{TODO}}$ ordered with the reference relation $\prec_s$ and remove those equivalence classes from $\mathcal{C}_{\texttt{TODO}}$. $\mathcal{C}_{\texttt{all}}$ contains all equivalence classes that are stepwise aligned with a minimal representative selection function $r$. In each step, we also align axioms $\mathcal{T}_{\texttt{out}}$ corresponding to the partition $\mathcal{T}^0$ with $r$ by replacing concepts with the representative $r([C])$ fixed in Line 6.

In Line 11, we build partitions for non-singleton equivalence classes. In Lines 12-14, we compute the non-redundant part of $\mathcal{T}_{\texttt{out}}$. The function $\texttt{compose}(\mathcal{T}_{\texttt{out}})$ in Line 16 composes subsumption axioms with identical left-hand sides into a single axiom.

Clearly, Algorithm 1 runs in PTIME. Note that the algorithm remains tractable only assuming the tractability of reasoning in the underlying logic. Otherwise, the complexity of reasoning dominates. In principle, the result could be obtained after computing the representatives for each equivalence class by simply selecting all subsumption relations between classes. However, this would result in a less efficient implementation with large intermediary results.

**Theorem 4.** *Let $\mathcal{T}$ be an acyclic $\mathcal{EL}$ TBox. Algorithm 1 computes a minimal equivalent TBox in* PTIME.

**Proof Sketch.** The minimality of the resulting TBox follows from Theorem 3 if we show that the computed representative selection function $r$ is indeed minimal and valid. This can be shown by induction on the structure of our reference relations. By assumption, this structure is finite and acyclic. The basis of induction is the case that all equivalence classes that do not reference any other equivalence classes. For such equivalence classes it holds that $C^* = C$, i.e., all class elements are known and a minimal element can effectively be chosen. Additionally, it is impossible to choose a value that would make $r$ invalid. Thus, $r$ is minimal and valid. This reflects the initial computation state, where we choose values of $r$ for equivalence classes that do not reference any other equivalence classes. At a later stage we choose a value for an equivalence class $[C]$ after choosing the values of $r$ for all classes referenced from it. Assuming that the values of $r$ for all classes directly or indirectly referenced from $[C]$ are valid and minimal, also any minimal value in $[C]$ is valid and minimal, since $\mathcal{T}$ and all equivalence classes have been iteratively aligned with the already fixed values of $r$. □

## 6 Experimental Results

For our evaluation, we have implemented the algorithm using the latest version of OWL API and Hermit reasoner. We have used an optimized version of Algorithm 1, where entailment checking is done in two phases, the first of which can be run by several threads.

A selection of publicly available ontologies (as shown in Table 1) that vary in size and expressivity have been used in the experiments[5]. Table 2 shows the number $|\mathtt{CON}_o(\mathcal{T})|$ of occurrences of complex concepts $\mathtt{CON}(\mathcal{T}) = \mathtt{sub}(\mathcal{T}) \setminus \mathrm{sig}_C(\mathcal{T})$ in the first two columns (the original value followed by the new value relative to the original one). The two subsequent columns show the number of pairwise different complex concepts $|\mathtt{CON}(\mathcal{T})|$. The last two columns show $\int(\mathcal{T})$ – the size of each ontology measured as the number of occurrences of entities in $\mathrm{sig}(\mathcal{T})$.

The implementation was first applied to Snomed [10]. However, the available fully-fledged reasoners Pellet and Hermit run out of heap space when classifying the ontology even with 10 GB memory assigned to the corresponding Java process. The ELK reasoner [11] is capable of classifying Snomed, but it does not currently implement entailment, which is essential for our implementation.

From the ontologies used in our experiments, only Snomed did not satisfy the acyclicity conditions for $\prec_s$ sufficient to guarantee termination of our algorithm. On the

[5] The wine ontology can be retrieved from http://www.w3.org/TR/2003/ PR-owl-guide-20031209/wine. All other ontologies used can be found in the TONES ontology repository at http://owl.cs.manchester.ac.uk/repository

| | $\|\mathcal{T}\|$ | $\int(\mathcal{T})/\|\mathcal{T}\|$ | $\mathtt{CON}(\mathcal{T})/\|\mathcal{T}\|$ | $\mathtt{CON}_o(\mathcal{T})/\|\mathcal{T}\|$ | Logic |
|---|---|---|---|---|---|
| Snomed | 83,259 | 4.99 | 1.14 | 2.57 | $\mathcal{EL}++$ |
| Gene Ontology | 42656 | 3.37 | 1.20 | 0.27 | $\mathcal{EL}++$ |
| NCI | 97811 | 1.10 | 0.00 | 0.14 | $\mathcal{ALCH}(\mathcal{D})$ |
| Galen | 4735 | 2.81 | 0.52 | 1.13 | $\mathcal{ALEHIF}+$ |
| Adult Mouse | 3464 | 2.48 | 0.15 | 0.48 | $\mathcal{EL}++$ |
| Wine | 657 | 1.03 | 0.21 | 0.40 | $\mathcal{SHOIN}(\mathcal{D})$ |
| Nautilus | 38 | 2.18 | 0.29 | 0.40 | $\mathcal{ALCHF}(\mathcal{D})$ |
| Cell | 1264 | 2.16 | 0.09 | 0.16 | $\mathcal{EL}++$ |
| DOLCE-lite | 351 | 1.42 | 0.13 | 0.14 | $\mathcal{SHIF}$ |
| Software | 238 | 25.21 | 2.60 | 7.26 | $\mathcal{ALHN}(\mathcal{D})$ |
| Family Tree | 36 | 6.19 | 1.02 | 1.33 | $\mathcal{SHIN}(\mathcal{D})$ |
| General Ontology | 8803 | 0.48 | 0.03 | 0.03 | $\mathcal{ALCHOIN}(\mathcal{D})$ |
| Substance | 609 | 2.33 | 0.22 | 0.36 | $\mathcal{ALCHO}(\mathcal{D})$ |
| Generations | 38 | 1.87 | 0.58 | 1.21 | $\mathcal{ALCOIF}$ |
| Periodic Table | 58 | 1.38 | 0.38 | 0.43 | $\mathcal{ALU}$ |

**Table 1.** Properties of ontologies used in experiments.

one hand, Snomed contains cyclic concept definitions. For instance, `Mast_cell_leukemia` is defined by means of the corresponding equivalence axiom as

$$\begin{aligned}
&\mathtt{Leukemia_disease} \ \sqcap \\
&\mathtt{Mast_cell_malignancy} \ \sqcap \\
&\exists \mathtt{RoleGroup}. \\
&\quad (\exists \mathtt{Associated_morphology}.\mathtt{Mast_cell_leukemia} \ \sqcap \\
&\quad \exists \mathtt{Finding_site}.\mathtt{Hematopoietic_system_structure})) \ \sqcap \\
&\exists \mathtt{RoleGroup}.( \\
&\quad \exists \mathtt{Has_definitional_manifestation}.\mathtt{White_blood_cell_finding})
\end{aligned}$$

On the other hand, Snomed contains a cyclic reference relation between the concepts `Wound` and `Wound_finding`, which is the only cyclic dependency with more than one element.

We have manually evaluated how the rewriting has affected ontologies. In all cases where concepts became smaller, the improvement has been achieved by either elimination of redundant axioms or exchanging complex expressions by atomic concepts.

In case of the Galen ontology [1], the algorithm managed to reduce the number of occurrences of complex concepts by 955, which is 17%. The size of the ontology in number of references was reduced by 1130 (9%). The number of distinct complex concepts used in the ontology was reduced by 76 (3%). The situation is similar for the NCI [12] ontology.

The other large medical ontology – Gene Ontology [13] – does not contain any equivalent concepts, i.e., each equivalence class has only one element. The current algorithm did not find any possibility to rewrite the ontology. The same holds for Adult Mouse and Periodic Table ontologies.

| | $\texttt{CON}_o(\mathcal{T})$ | | $\lvert\texttt{CON}(\mathcal{T})\rvert$ | | $\int(\mathcal{T})$ | |
|---|---|---|---|---|---|---|
| Snomed | 213,856 | – | 95,315 | – | 415,541 | – |
| Gene Ontology | 11,686 | 1 | 8,508 | 1 | 143,900 | 1 |
| NCI | 13,961 | 0.87 | 4,000 | 0.99 | 107,841 | 0.94 |
| Galen | 5,368 | 0.83 | 2,475 | 0.97 | 13,285 | 0.91 |
| Adult Mouse | 1,649 | 0.99 | 507 | 1 | 8,575 | 0.99 |
| Wine | 262 | 0.89 | 141 | 0.98 | 677 | 0.93 |
| Nautilus | 15 | 1 | 11 | 1 | 83 | 0.86 |
| Cell | 206 | 0.87 | 114 | 0.96 | 2,732 | 0.96 |
| DOLCE-lite | 49 | 0.92 | 46 | 0.98 | 497 | 0.66 |
| Software | 1,728 | 0.81 | 620 | 1 | 6,001 | 0.81 |
| Family Tree | 48 | 0.77 | 37 | 0.78 | 223 | 0.83 |
| General Ontology | 281 | 0.83 | 278 | 0.83 | 4,182 | 0.83 |
| Substance | 221 | 1 | 135 | 1 | 1,417 | 0.95 |
| Generations | 46 | 0.65 | 22 | 1 | 71 | 0.90 |
| Periodic Table | 25 | 1 | 22 | 1 | 80 | 1 |

**Table 2.** Minimization results for different ontologies.

Results for the other, relatively small ontologies are similar to those of Galen and in some cases more prominent (Table 2). The highest improvement (66% of $\int(\mathcal{T})$) was achieved in the DOLCE-Lite ontology [14].

## 7 Related Work

The work on knowledge compilation [15] is closely related to the work presented in this paper. Knowledge compilation is a family of approaches, in which a knowledge base is transformed in an off-line phase into a normal form, for which reasoning is cheaper. The hope is that the one-off cost of the initial preprocessing will be justified by the computational savings made on subsequent reasoning. One of such normal forms proposed in description logics is the prime implicates normal form for $\mathcal{ALC}$ ontologies [4]. Prime implicates of a logical formula are defined to be their strongest clausal consequences. Concepts in the prime implicates normal form are expected to be easier to read and understand. Reasoning is also expected to be more efficient for knowledge bases in this normal form. For example, concept subsumption can be tested in quadratic time. However, the problem with such normal forms is the blowup caused by the transformation. For $\mathcal{ALC}$ ontologies, a doubly-exponential blowup in the concept size can occur. Given that reasoning in $\mathcal{ALC}$ is PSPACE-complete [16], such a transformation can be disadvantageous in general.

Grimm et al. [3] propose two different algorithms for eliminating semantically redundant axioms from ontologies, which is one of the sources of non-succinctness. However, as shown in Section 3, it does not guarantee that we obtain a minimal TBox in $[(]\mathcal{T})$. The advantage of this restricted approach to improving succinctness is that the result contains only axioms that are familiar to the users of the ontology.

Work on laconic and precise justifications [17] (minimal parts of the ontology implying a particular axiom or axioms) is also related to this paper. The authors propose an algorithm for computing laconic justifications – justifications that do not contain any logically superfluous parts. Laconic justifications can then be used to derive precise justifications – justifications that consist of flat, small axioms, and are important for the generation of semantically minimal repairs.

Nikitina et al. [18] propose an algorithm for an efficient handling of redundancy in inconsistent ontologies during their repair. Similarly to the approach by Grimm et al. axioms are considered as atoms that cannot be further separated into parts.

## 8 Summary and Outlook

We have considered the problem of finding minimal equivalent representations for ontologies expressed in the lightweight description logic $\mathcal{EL}$. We have shown that the task of finding such a representation (or rather: its related decision problem) is NP-complete. Further, we have identified a class of TBoxes for which the problem is tractable. We have implemented a polynomial algorithm for minimizing the above class of TBoxes. For general TBoxes, the algorithm can be used as a heuristic. We have implemented the algorithm and presented experimental results, which show that the complexity of various existing ontologies can be improved. For instance, in case of Galen, the number of complex concepts occurrences could be reduced by 955 and the number of references to atomic concepts and roles by 1130.

There are various natural extensions of this work. Inspired by recent results on uniform interpolation in $\mathcal{EL}$ [8], the problem can be extended to finding minimal representations for ontologies using a signature extension. The results in [8] imply that, even for the minimal equivalent representation of an ontology, an up to triple-exponentially more succinct representation can be obtained by extending its signature. Auxiliary concept symbols are therefore important contributors towards the succinctness of ontologies, e.g., used as shortcuts for complex $\mathcal{EL}$ concepts or disjunctions thereof. The results of our evaluation indicate that there are many complex concept expression that occur repeatedly in ontologies but do not have an equivalent atomic concept that could be used instead. Therefore, introducing names for such frequently used concepts could yield a further decrease of the ontology's complexity.

The results obtained within this paper can be transferred to the context of ontology reuse, where rewriting is applied to obtain a compact representation of the facts about a subset of terms [19], in particular in its extended form as suggested above.

Finally, minimizing representations is an interesting problem for knowledge representation formalisms in general, and similar questions can (and should) be asked for more expressive ontology languages.

## References

1. Rector, A., Gangemi, A., Galeazzi, E., Glowinski, A., Rossi-Mori, A.: The GALEN CORE model schemata for anatomy: Towards a re-usable application-independent model of medical concepts. In: Proceedings of the 12th International Congress of the European Federation for Medical Informatics (MIE 1994). (1994) 229–233

2. Baader, F., Calvanese, D., McGuinness, D., Nardi, D., Patel-Schneider, P.: The Description Logic Handbook: Theory, Implementation and Applications. Cambridge University Press (2003)
3. Grimm, S., Wissmann, J.: Elimination of redundancy in ontologies. In: Proceedings of the 8th Extended Semantic Web Conference (ESWC 2011). (2011) 260–274
4. Bienvenu, M.: Prime implicates and prime implicants: From propositional to modal logic. Journal of Artificial Intelligence Research (JAIR) **36** (2009) 71–128
5. Motik, B., Cuenca Grau, B., Horrocks, I., Wu, Z., Fokoue, A., Lutz, C., eds.: OWL 2 Web Ontology Language: Profiles. W3C Recommendation (27 October 2009) Available at http://www.w3.org/TR/owl2-profiles/.
6. OWL Working Group, W.: OWL 2 Web Ontology Language: Document Overview. W3C Recommendation (27 October 2009) Available at http://www.w3.org/TR/owl2-overview/.
7. Baader, F., Brandt, S., Lutz, C.: Pushing the $\mathcal{EL}$ envelope. In: Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI 2005). (2005) 364–369
8. Nikitina, N., Rudolph, S.: ExpExpExplosion: Uniform interpolation in general EL terminologies. In: Proceedings of the 20th European Conference on Artificial Intelligence (ECAI 2012). (2012) 618–623 (Shortlisted for best paper awards).
9. Konev, B., Walther, D., Wolter, F.: Forgetting and uniform interpolation in large-scale description logic terminologies. In: Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI 2009). (2009) 830–835
10. Spackman, K.A., Campbell, K.E., Cote, R.A.: Snomed rt: A reference terminology for health care. In: Proceedings of the AIMA Fall Symposium. (1997) 640–644
11. Kazakov, Y., Krötzsch, M., Simančík, F.: ELK reasoner: Architecture and evaluation. In: Proceedings of the OWL Reasoner Evaluation Workshop 2012 (ORE'12). Volume 858 of CEUR Workshop Proceedings., CEUR-WS.org (2012)
12. Sioutos, N., Coronado, S.d., Haber, M.W., Hartel, F.W., Shaiu, W.L., Wright, L.W.: Nci thesaurus: A semantic model integrating cancer-related clinical and molecular information. Journal of Biomedical Informatics **40**(1) (February 2007) pp.30–43
13. Ashburner, M.: Gene ontology: Tool for the unification of biology. Nature Genetics **25** (2000) pp.25–29
14. Gangemi, A., Guarino, N., Masolo, C., Oltramari, A., Schneider, L.: Sweetening ontologies with dolce. In: Proceedings of the 13th International Conference on Knowledge Engineering and Knowledge Management. Ontologies and the Semantic Web. (2002) 166–181
15. Darwiche, A., Marquis, P.: A knowledge compilation map. Journal of Artificial Intelligence Research (JAIR) **17** (2002) 229–264
16. Schmidt-Schauß, M., Smolka, G.: Attributive concept descriptions with complements. Artificial Intelligence **48**(1) (1991) 1–26
17. Horridge, M., Parsia, B., Sattler, U.: Laconic and precise justifications in owl. In: Proceedings of the 7th International Semantic Web Conference (ISWC 2008). (2008) 323–338
18. Nikitina, N., Rudolph, S., Glimm, B.: Reasoning-supported interactive revision of knowledge bases. In: Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI 2011). (2011) 1027–1032
19. Nikitina, N., Glimm, B.: Hitting the sweetspot: Economic rewriting of knowledge bases. In: Proceedings of the 11th International Semantic Web Conference (ISWC 2012). (2012) 394–409