

Learning Description Logic Concepts: When can Positive and Negative Examples be Separated?

Maurice Funk¹, Jean Christoph Jung¹, Carsten Lutz¹, Hadrien Pulcini² and Frank Wolter²

¹Department of Computer Science, University of Bremen

²Department of Computer Science, University of Liverpool

Abstract

Learning description logic (DL) concepts from positive and negative examples given in the form of labeled data items in a KB has received significant attention in the literature. We study the fundamental question of when a separating DL concept exists and provide useful model-theoretic characterizations as well as complexity results for the associated decision problem. For expressive DLs such as \mathcal{ALC} and \mathcal{ALCQL} , our characterizations show a surprising link to the evaluation of ontology-mediated conjunctive queries. We exploit this to determine the combined complexity (between EXPTIME and NEXPTIME) and data complexity (second level of the polynomial hierarchy) of separability. For the Horn DL \mathcal{EL} , separability is EXPTIME-complete both in combined and in data complexity while for its modest extension \mathcal{ELI} it is even undecidable. Separability is also undecidable when the KB is formulated in \mathcal{ALC} and the separating concept is required to be in \mathcal{EL} or \mathcal{ELI} .

1 Introduction

An important challenge for adopting ontologies in practical applications is the knowledge acquisition bottleneck, that is, the significant time and effort it takes to build the required ontologies. As a promising approach to help overcoming this difficulty, the varied field of *ontology learning* has received a lot of attention in the last two decades, see [Lehmann and Völker, 2014] for a recent overview. A prominent line of research within ontology learning is *concept learning*, also called *concept induction*, where the aim is to learn a structured class description (a *concept*) formulated in a relevant ontology language from positive and negative examples, given an already available ontology that contains background knowledge [Lehmann *et al.*, 2014]. Applications of concept learning include the support of ontology development and the construction of concept based classifiers [Bühmann *et al.*, 2016; Sarker *et al.*, 2017].

Since description logics (DLs) form the basis for the OWL family of ontology languages, DL concept learning has received particularly much attention [Lehmann and Hitzler, 2010; Lisi, 2012; Tran *et al.*, 2014; Fanizzi *et al.*, 2018]. The

precise formulation is as follows. Given a knowledge base (KB) $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ where \mathcal{T} is a TBox that is formulated in a DL \mathcal{L}_T and serves as an ontology providing background knowledge and \mathcal{A} is an ABox in which individuals from designated sets P and N serve as positive and negative examples, respectively, find a concept C formulated in a DL \mathcal{L}_S that *separates* the positive from the negative examples, that is, $\mathcal{K} \models C(a)$ for all $a \in P$ and $\mathcal{K} \not\models C(a)$ for all $a \in N$. In addition to separation, one also wants to achieve that the learned concept C generalizes the positive examples in a meaningful way, classifying new examples accordingly.

As a prominent system for DL concept learning, we mention DL LEARNER. It encompasses several learning algorithms that support a range of DLs, including expressive ones such as \mathcal{ALC} and \mathcal{ALCQ} , Horn DLs such as \mathcal{EL} , and even full OWL 2 [Bühmann *et al.*, 2018; Bühmann *et al.*, 2016]. Like competing systems such as DL-FOIL, YINYANG, and PFOIL-DL [Fanizzi *et al.*, 2018; Iannone *et al.*, 2007; Straccia and Mucci, 2015], DL LEARNER uses a carefully crafted *refinement operator* along with various heuristics to learn concepts that provide an as good as possible generalization of the given examples, avoiding overfitting. The study of such operators originated in [Badea and Nienhuys-Cheng, 2000], see also [Lehmann and Haase, 2009; Lehmann and Hitzler, 2010]. If possible, refinement operators are designed so that the resulting algorithm terminates on any input and is complete in the sense that whenever there is a concept that separates the positive and negative examples in the input, then such a concept is indeed learned.

The aim of this paper is to investigate the fundamental question of when a separating concept exists for a *learning instance* (\mathcal{K}, P, N) , considering the most popular choices of DLs for the TBox language \mathcal{L}_T and the separation language \mathcal{L}_S . Our main contributions are model-theoretic characterizations that give important insight into when this is the case and a precise analysis of the computational complexity of separability viewed as a decision problem, which we refer to as $(\mathcal{L}_T, \mathcal{L}_S)$ *concept separability* and as \mathcal{L} *concept separability* when $\mathcal{L}_T = \mathcal{L}_S = \mathcal{L}$. We also consider *concept definability*, the special case of concept separability in which $P \cup N$ comprises all individuals from \mathcal{A} . In fact, all our complexity results hold for both separability and definability.

We believe that these problems are relevant both from a practical and from a theoretical perspective. In fact, complex-

ity lower bounds for concept separability point to an inherent complexity that no practical system that aims for completeness can avoid. Undecidability results even mean that there can be no practical learning system that is both terminating and complete. From the viewpoint of machine learning theory, a separating concept corresponds to a consistent hypothesis and understanding the existence of such a hypothesis is considered the most fundamental problem for exploring the version space which in our setup is the space of all separating concepts [Hirsh *et al.*, 2004]. The associated decision problem is often called the *consistency problem*, and it is known to be closely related to PAC learnability [Pitt and Valiant, 1988; Kietz, 1993].

We cover the expressive DLs \mathcal{ALC} , \mathcal{ALCI} , \mathcal{ALCQ} , and \mathcal{ALCQI} as well as the Horn DLs \mathcal{EL} and \mathcal{ELI} . For the former, overfitting is a risk because the disjunction operator available in such DLs enables the construction of separating concepts that do not provide the desired generalization of the positive examples. Nevertheless, most practical systems such as DL LEARNER work with expressive DLs and avoid overfitting by using appropriate refinement operators. Horn DLs do not admit disjunction and therefore are not prone to overfitting. On the other hand, they provide less separating power and, as we show, tend to incur higher computational (worst-case) cost for learning.

For expressive DLs, we start with initial characterizations in terms of (some form of) bisimulations and then improve them to more refined characterizations based on homomorphisms. Interestingly and unexpectedly to us, these establish a tight link between concept separability and the evaluation of ontology-mediated queries (OMQs) based on unions of rooted conjunctive queries [Calvanese *et al.*, 2013; Bienvenu *et al.*, 2014]. Here, ‘rooted’ means that queries have at least one answer variable and are connected. We use the OMQ connection to obtain complexity upper and lower bounds. In fact, \mathcal{L} concept separability is NEXPTIME-complete for $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCI}, \mathcal{ALCQ}\}$ while \mathcal{ALCQI} concept separability is only EXPTIME-complete. This refers to combined complexity where all components of the learning instance are part of the input. We also study data complexity where the ABox is the only input while the TBox is fixed and thus of constant size. In all expressive DLs mentioned above, concept separability is Σ_2^P -complete in data complexity.

The connection to OMQ evaluation does not extend to Horn DLs, for which we use characterizations based on products of universal models and simulations. Based on these, we show that $(\mathcal{L}_T, \mathcal{EL})$ concept separability is EXPTIME-complete for $\mathcal{L}_T \in \{\mathcal{EL}, \mathcal{ELI}\}$, both in combined complexity and in data complexity. Rather surprisingly, we also prove that \mathcal{ELI} concept separability is undecidable, thus ruling out terminating and complete learning systems. The proof is by a subtle reduction of a tiling problem. We finally consider $(\mathcal{L}_T, \mathcal{L}_S)$ concept separability where \mathcal{L}_T is any of the expressive DLs mentioned above and \mathcal{L}_S is \mathcal{EL} or \mathcal{ELI} . These problems also turn out to be undecidable.

A stronger version of concept separability that is also considered in the literature requires that $\mathcal{K} \models \neg C(a)$ for all $a \in N$, rather than only $\mathcal{K} \not\models C(a)$. Clearly, this is a meaningful notion only for DLs with negation. We present some

first results also for this version of concept separability, giving model-theoretic characterizations for expressive DLs based on bisimulations and proving that for $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCI}\}$, \mathcal{L} concept separability is EXPTIME-complete in combined complexity and CONP-complete in data complexity.

The long version with appendix is available at <http://www.informatik.uni-bremen.de/tdki/research/papers.html>.

2 Preliminaries

We introduce the basics of DLs as required for this paper, for full details see [Baader *et al.*, 2017]. Let N_C be a set of *concept names* and N_R a set of *role names*, both countably infinite. A *role* is either a role name or an *inverse role* r^- , r a role name. For uniformity, we identify $(r^-)^-$ with r . An \mathcal{ALCQI} *concept* is formed according to the syntax rule

$$C, D ::= \top \mid A \mid \neg C \mid C \sqcap D \mid (\geq n r C)$$

where A ranges over concept names, r over roles, and n over \mathbb{N} . We use $\exists r.C$ as an abbreviation for $(\geq 1 r C)$, $C \sqcup D$ for $\neg(\neg C \sqcap \neg D)$, $\forall r.C$ for $\neg\exists r.\neg C$, and $(\leq n r C)$ for $\neg(\geq n + 1 r C)$.

There are several fragments of \mathcal{ALCQI} that are relevant for this paper. An \mathcal{EL} *concept* admits as constructors only the top concept ‘ \top ’, conjunction ‘ \sqcap ’, and existential restriction ‘ $\exists r.C$ ’ with r a role name (but not an inverse role). \mathcal{ALC} *concepts* additionally admit negation ‘ \neg ’. Further constructors are indicated by concatenation of a corresponding letter where Q stands for at least restrictions ‘ $(\geq n r C)$ ’ and I for inverse roles. This explains the name \mathcal{ALCQI} and allows us to refer to fragments such as \mathcal{ELI} and \mathcal{ALCI} . We refer to \mathcal{ALC} and extensions thereof as *expressive DLs* and to \mathcal{EL} and \mathcal{ELI} as *Horn DLs*.

For any of the DLs \mathcal{L} introduced above, an \mathcal{L} *TBox* is a finite set of *concept inclusions* (CIs) $C \sqsubseteq D$, where C and D are \mathcal{L} concepts. Let N_I be a countably infinite set of *individual names*. An *ABox* \mathcal{A} is a finite set of *concept assertions* $A(a)$ and *role assertions* $r(a, b)$ where $A \in N_C$, $r \in N_R$, and $a, b \in N_I$. We use $\text{ind}(\mathcal{A})$ to denote the set of all individual names that occur in \mathcal{A} . An ABox is associated with a directed graph $G_{\mathcal{A}} = (\text{ind}(\mathcal{A}), \{(a, b) \mid r(a, b) \in \mathcal{A}\})$. We use graph theoretic terminology when speaking about ABoxes, implicitly referring to the graph $G_{\mathcal{A}}$, speaking for example about a being reachable from b for $a, b \in \text{ind}(\mathcal{A})$. In technical constructions, we will sometimes use *extended ABoxes* in which concept assertions take the more general form $C(a)$, C a concept of the DL under consideration. An \mathcal{L} *knowledge base* (KB) $(\mathcal{T}, \mathcal{A})$ consists of an \mathcal{L} TBox \mathcal{T} and an ABox \mathcal{A} , and *extended KBs* admit extended ABoxes. We use $\text{sub}(\mathcal{K})$ to denote the set of subconcepts of concepts that occur in \mathcal{K} .

As usual, the semantics of DLs is defined in terms of *interpretations* $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, where $\Delta^{\mathcal{I}}$ is a non-empty set and $\cdot^{\mathcal{I}}$ maps each concept name $A \in N_C$ to a subset $A^{\mathcal{I}}$ of $\Delta^{\mathcal{I}}$ and each role name $r \in N_R$ to a binary relation $r^{\mathcal{I}}$ on $\Delta^{\mathcal{I}}$. We refer to [Baader *et al.*, 2017] for details on how to extend $\cdot^{\mathcal{I}}$ to compound concepts. An interpretation \mathcal{I} *satisfies* a CI $C \sqsubseteq D$ if $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$ and is a *model* of a TBox \mathcal{T} if it satisfies all inclusions in \mathcal{T} . An interpretation is a *model* of an ABox \mathcal{A} if it *satisfies* all assertions in \mathcal{A} , that is, $a \in A^{\mathcal{I}}$ if

[Atom]	for all $(d, e) \in S$: $d \in A^{\mathcal{I}}$ iff $e \in A^{\mathcal{J}}$
[AtomR]	if $(d, e) \in S$ and $d \in A^{\mathcal{I}}$, then $e \in A^{\mathcal{J}}$
[Forth]	if $(d, e) \in S$ and $d' \in \text{succ}_r^{\mathcal{I}}(d)$, then there is a $e' \in \text{succ}_r^{\mathcal{J}}(e)$ with $(d', e') \in S$.
[Back]	dual of [Forth]
[QForth]	if $(d, e) \in S$ and $D \subseteq \text{succ}_r^{\mathcal{I}}(d)$ finite, then there is a $E \subseteq \text{succ}_r^{\mathcal{J}}(e)$ such that S contains a bijection between D and E .
[QBack]	dual of [QForth]

 Figure 1: Conditions on $S \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{J}}$.

$A(a) \in \mathcal{A}$ and $(a, b) \in r^{\mathcal{I}}$ if $r(a, b) \in \mathcal{A}$. We thus make the standard names assumption. An interpretation is a *model* of a KB \mathcal{K} if it is a common model of its TBox and ABox and \mathcal{K} is *satisfiable* if it has a model. The assertion $C(a)$ is a *consequence of the KB \mathcal{K}* , in symbols $\mathcal{K} \models C(a)$, if $a \in C^{\mathcal{I}}$ for all models \mathcal{I} of \mathcal{K} .

We next recall model-theoretic characterizations of when elements in interpretations are indistinguishable by concepts formulated in one of the DLs \mathcal{L} introduced above. A *pointed interpretation* is a pair \mathcal{I}, d with \mathcal{I} an interpretation and $d \in \Delta^{\mathcal{I}}$. For pointed interpretations \mathcal{I}, d and \mathcal{J}, e , we write $\mathcal{I}, d \equiv_{\mathcal{L}} \mathcal{J}, e$ and say that \mathcal{I}, d and \mathcal{J}, e are *\mathcal{L} -equivalent* if $d \in C^{\mathcal{I}}$ iff $e \in C^{\mathcal{J}}$ for all \mathcal{L} concepts C .

As for the model-theoretic characterizations, we start with \mathcal{ALC} . A relation $S \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{J}}$ is a *bisimulation* if conditions [Atom], [Forth] and [Back] from Figure 1 hold, where A ranges over all concept names, r over all role names, $\text{succ}_r^{\mathcal{I}}(d) = \{d' \in \Delta^{\mathcal{I}} \mid (d, d') \in r^{\mathcal{I}}\}$, and ‘dual’ refers to swapping \mathcal{I}, d, d' and \mathcal{J}, e, e' . We write $\mathcal{I}, d \sim_{\mathcal{ALC}} \mathcal{J}, e$ and call \mathcal{I}, d and \mathcal{J}, e *bisimilar* if there exists a bisimulation S such that $(d, e) \in S$. For \mathcal{ALCQ} , we define $\sim_{\mathcal{ALCQ}}$ by replacing bisimulations with *counting bisimulations*, defined as bisimulations, but with [Forth] and [Back] replaced by [QForth] and [QBack]. For \mathcal{ALCI} and \mathcal{ALCQI} , we define $\sim_{\mathcal{ALCI}}$ and $\sim_{\mathcal{ALCQI}}$ analogously, but now demanding that in all conditions in Figure 1, r additionally ranges over inverse roles.

For $\mathcal{L} \in \{\mathcal{EL}, \mathcal{ELI}\}$, rather than characterizing $\equiv_{\mathcal{L}}$ we consider the non-symmetric relation $\mathcal{I}, d \leq_{\mathcal{L}} \mathcal{J}, e$ which holds if $d \in C^{\mathcal{I}}$ implies $e \in C^{\mathcal{J}}$ for all \mathcal{L} concepts C . A relation $S \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{J}}$ is an *\mathcal{EL} simulation* from \mathcal{I} to \mathcal{J} if it satisfies [AtomR] and [Forth] from Figure 1 where again A ranges over all concept names and r over all role names. *\mathcal{ELI} simulations* are defined in the same way, but with r ranging also over inverse roles. We write $\mathcal{I}, d \preceq_{\mathcal{L}} \mathcal{J}, e$ if there exists an \mathcal{L} simulation S from \mathcal{I} to \mathcal{J} with $(d, e) \in S$.

The next lemma summarizes the model-theoretic characterizations for all relevant DLs [Lutz *et al.*, 2011; Goranko and Otto, 2007]. An interpretation \mathcal{I} has *finite outdegree* if the directed graph $G_{\mathcal{I}} = (\Delta^{\mathcal{I}}, \bigcup_{r \in \mathcal{N}_R} r^{\mathcal{I}})$ has.

Lemma 1 *Let \mathcal{I}, d and \mathcal{J}, e be pointed interpretations and let \mathcal{J} have finite outdegree. Then*

1. for $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCI}, \mathcal{ALCQ}, \mathcal{ALCQI}\}$, $\mathcal{I}, d \equiv_{\mathcal{L}} \mathcal{J}, e$ iff $\mathcal{I}, d \sim_{\mathcal{L}} \mathcal{J}, e$;
2. for $\mathcal{L} \in \{\mathcal{EL}, \mathcal{ELI}\}$, $\mathcal{I}, d \leq_{\mathcal{L}} \mathcal{J}, e$ iff $\mathcal{I}, d \preceq_{\mathcal{L}} \mathcal{J}, e$.

The ‘if’ directions also hold if \mathcal{J} is of infinite outdegree.

3 Separability and Definability

We introduce concept separability and concept definability, the main notions studied in this paper. We also provide illustrating examples and give initial model-theoretic characterizations in all relevant DLs.

Definition 1 *Let $\mathcal{L}_T, \mathcal{L}_S$ be DLs. An \mathcal{L}_T learning instance is a triple (\mathcal{K}, P, N) with $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ an \mathcal{L}_T KB and $P, N \subseteq \text{ind}(\mathcal{A})$ non-empty sets of positive and negative examples. An \mathcal{L}_S solution to (\mathcal{K}, P, N) is an \mathcal{L}_S concept C such that*

1. $\mathcal{K} \models C(a)$ for all $a \in P$ and
2. $\mathcal{K} \not\models C(a)$ for all $a \in N$.

Any TBox language \mathcal{L}_T and separation language \mathcal{L}_S give rise to a decision problem of *concept separability*.

PROBLEM: $(\mathcal{L}_T, \mathcal{L}_S)$ concept separability
 INPUT: \mathcal{L}_T learning instance (\mathcal{K}, P, N)
 QUESTION: Does (\mathcal{K}, P, N) have an \mathcal{L}_S solution?

We speak of \mathcal{L} *concept separability* when $\mathcal{L}_T = \mathcal{L}_S = \mathcal{L}$. We also study $(\mathcal{L}_T, \mathcal{L}_S)$ *concept definability*, the special case of $(\mathcal{L}_T, \mathcal{L}_S)$ concept separability in which P covers *all* positive examples in the KB, that is, inputs are learning instances (\mathcal{K}, P, N) such that $N = \text{ind}(\mathcal{A}) \setminus P$. We now give illustrating examples. When we claim that there is no solution, this is a direct consequence of the characterizations provided later. In the following examples and also later in this paper, we assume that $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ unless explicitly stated otherwise.

Example 1 (1) Let (\mathcal{K}, P, N) be defined by $\mathcal{T} = \emptyset$,

$$\mathcal{A} = \{\text{authorOf}(a_i, c_i) \mid i = 1, 2, 3\} \cup \{\text{IJCAIpub}(c_1), \text{AIJpub}(c_2), \text{GraphicNovel}(c_3)\},$$

$P = \{a_1, a_2\}$, and $N = \{a_3\}$. Then $\exists \text{authorOf}.\text{IJCAIpub} \sqcup \text{AIJpub}$ is an \mathcal{ALC} solution, but there is neither an \mathcal{EL} solution nor an \mathcal{ELI} solution.

(2) Let (\mathcal{K}, P, N) be as in (1), but with \mathcal{T} replaced by the \mathcal{EL} TBox

$$\mathcal{T} = \{\exists \text{authorOf}.\text{IJCAIpub} \sqsubseteq \text{AIResearcher}, \exists \text{authorOf}.\text{AIJpub} \sqsubseteq \text{AIResearcher}\}.$$

Then AIResearcher is an \mathcal{EL} solution.

(3) Let (\mathcal{K}, P, N) be defined by $\mathcal{T} = \emptyset$,

$$\mathcal{A} = \{\text{authorOf}(a, c_1), \text{authorOf}(a, c_2), \text{authorOf}(b, c_3)\},$$

$P = \{a\}$, and $N = \{b\}$. Due to the standard names assumption, c_1 and c_2 are distinct objects and thus $(\geq 2 \text{ authorOf } \top)$ is an \mathcal{ALCQ} solution. However, there is no \mathcal{ALCI} solution, even when \mathcal{T} is replaced with *any* \mathcal{ALCI} TBox.

(4) Let (\mathcal{K}, P, N) be as in (3), but with P and N swapped. There is no \mathcal{ALCQI} solution even when \mathcal{T} is replaced with *any* \mathcal{ALCQI} TBox. Note that $(\leq 1 \text{ authorOf } \top)$ is not a solution as the semantics does not rule out additional authorOf successors in concrete models.

Observe that if \mathcal{L}_S is closed under conjunction, then $(\mathcal{L}_T, \mathcal{L}_S)$ concept separability can be reduced in polynomial time to the special case of $(\mathcal{L}_T, \mathcal{L}_S)$ concept separability that admits only input instances (\mathcal{K}, P, N) in which N is a

singleton. Indeed, if C is a solution to a learning instance (\mathcal{K}, P, N) , then it is a solution to $(\mathcal{K}, P, \{a\})$ for all $a \in N$. If, conversely, C_a is a solution to $(\mathcal{K}, P, \{a\})$ for all $a \in N$, then $\prod_{a \in N} C_a$ is a solution to (\mathcal{K}, P, N) . In what follows, we will thus mostly restrict our attention to this special case.

We study both the *combined complexity* of concept separability and definability and the *data complexity*, where only the ABox of the KB is regarded as the input to the decision problem while the TBox is fixed and thus of constant size. Note that the sizes of P and N are dominated by the size of the ABox. When not making explicit which complexity measure we speak about, we mean combined complexity.

We now give the announced characterizations of $(\mathcal{L}_S, \mathcal{L}_T)$ concept separability, starting with the case that \mathcal{L}_S is an expressive DL. The characterizations use the relations $\sim_{\mathcal{L}}$ from Lemma 1, \mathcal{L} being the separation language. The corresponding relation for the TBox language is not used.

Theorem 1 *Let $(\mathcal{K}, P, \{b\})$ be an \mathcal{ALCQI} learning instance. For $\mathcal{L}_S \in \{\mathcal{ALC}, \mathcal{ALCI}, \mathcal{ALCQ}, \mathcal{ALCQI}\}$, the following are equivalent:*

1. $(\mathcal{K}, P, \{b\})$ has an \mathcal{L}_S solution;
2. there exists a model \mathcal{I} of \mathcal{K} of finite outdegree such that for all $a \in P$ and all models \mathcal{J} of \mathcal{K} , $\mathcal{J}, a \not\sim_{\mathcal{L}_S} \mathcal{I}, b$.

The proof of Theorem 1 is rather standard, based on Lemma 1 and compactness.

We now turn to the case where \mathcal{L}_S is a Horn DL. The main difference to Theorem 1 is that a product construction is involved. Interestingly, our characterization is rather flexible regarding \mathcal{L}_T , which can even be \mathcal{ALCQI} . Let $\mathcal{I}_i, d_i, i \in I$, be a family of pointed interpretations. The (direct) product $\prod_{i \in I} \mathcal{I}_i, d$ is the pointed interpretation defined by

$$\begin{aligned} \Delta \prod \mathcal{I}_i &= \{d : I \rightarrow \bigcup_{i \in I} \Delta^{\mathcal{I}_i} \mid \text{for } i \in I : d(i) \in \Delta^{\mathcal{I}_i}\} \\ A \prod \mathcal{I}_i &= \{d \in \Delta \prod \mathcal{I}_i \mid \text{for } i \in I : d(i) \in A^{\mathcal{I}_i}\}, \quad A \in \mathbf{N}_C \\ r \prod \mathcal{I}_i &= \{(d, e) \mid \text{for } i \in I : (d(i), e(i)) \in r^{\mathcal{I}_i}\}, \quad r \in \mathbf{N}_R \end{aligned}$$

and $d(i) = d_i$, for all $i \in I$. We say that a set \mathcal{M} of models of a KB \mathcal{K} is \mathcal{L} complete if for every \mathcal{L} concept C and every $a \in \text{ind}(\mathcal{K})$, $\mathcal{K} \models C(a)$ iff $a \in C^{\mathcal{I}}$ for all $\mathcal{I} \in \mathcal{M}$. If, for example, \mathcal{K} is an \mathcal{ALCQI} KB, then the class of all ‘forest models’ of \mathcal{K} of finite outdegree is well-known to be \mathcal{ALCQI} complete.

Theorem 2 *Let $(\mathcal{K}, P, \{b\})$ be an \mathcal{ALCQI} learning instance. For $\mathcal{L}_S \in \{\mathcal{EL}, \mathcal{ELI}\}$ and \mathcal{M} a set of models of \mathcal{K} that is \mathcal{L}_S complete, the following are equivalent:*

1. $(\mathcal{K}, P, \{b\})$ has an \mathcal{L}_S solution;
2. there exists a model \mathcal{I} of \mathcal{K} of finite outdegree such that $\prod_{a \in P} (\prod_{\mathcal{J} \in \mathcal{M}} (\mathcal{J}, a)) \not\sim_{\mathcal{L}_S} \mathcal{I}, b$.

The proof uses Lemma 1 and the fact that for any \mathcal{ELI} concept C and family of pointed interpretations $(\mathcal{I}_i, d_i)_{i \in I}$, $d \in C^{\prod_{i \in I} \mathcal{I}_i}$ if and only if $d_i \in C^{\mathcal{I}_i}$ for all $i \in I$.

4 Characterizations for Expressive DLs

We establish characterizations of concept separability that are more refined than the initial one presented in Theorem 1. The

refined characterizations establish a surprising connection between concept separability and (U)CQ-evaluation on KBs. They also provide the foundation for decision procedures for concept separability that we develop later on. We start with \mathcal{ALCI} and then proceed via \mathcal{ALCQI} to \mathcal{ALC} and \mathcal{ALCQ} , for which the characterizations are slightly more technical.

A homomorphism h from an ABox \mathcal{A} to an interpretation \mathcal{I} is a mapping h from $\text{ind}(\mathcal{A})$ to $\Delta^{\mathcal{I}}$ such that $A(d) \in \mathcal{A}$ implies $h(d) \in A^{\mathcal{I}}$ and $r(c, d) \in \mathcal{A}$ implies $(h(c), h(d)) \in r^{\mathcal{I}}$. We write $\mathcal{A}, a \rightarrow \mathcal{I}, b$ if there exists a homomorphism h from \mathcal{A} to \mathcal{I} with $h(a) = b$.

Theorem 3 *Let $(\mathcal{K}, P, \{b\})$ be an \mathcal{ALCI} learning instance. Then the following conditions are equivalent:*

1. $(\mathcal{K}, P, \{b\})$ has an \mathcal{ALCI} solution;
2. there exists a model \mathcal{I} of \mathcal{K} such that for all $a \in P$, $\mathcal{A}, a \not\rightarrow \mathcal{I}, b$.

Theorem 3 is proved by showing the equivalence of Condition 2 in Theorems 1 and 3. One first observes that it suffices to consider models \mathcal{I} that contain no distinct d, e with $\mathcal{I}, d \sim_{\mathcal{ALCI}} \mathcal{I}, e$. Then the direction from Theorem 3 to Theorem 1 follows as the restriction of a bisimulation between \mathcal{J} and \mathcal{I} to the ABox individuals in \mathcal{A} is a homomorphism to \mathcal{I} . For the converse direction, one carefully constructs the required model \mathcal{J} of \mathcal{K} from \mathcal{I} using the given homomorphism.

To adapt Theorem 3 from \mathcal{ALCI} to \mathcal{ALCQI} , we use homomorphisms h from \mathcal{A} to \mathcal{I} that are *locally injective* for all roles r , that is, the restriction of h to the set $\text{succ}_r^{\mathcal{A}}(a) = \{b \mid r(a, b) \in \mathcal{A}\}$ is injective for all $a \in \text{ind}(\mathcal{A})$. Now $\mathcal{A}, a \rightarrow^i \mathcal{I}, b$ is defined like $\mathcal{A}, a \rightarrow \mathcal{I}, b$, but based on homomorphisms that are locally injective for all roles r .

Theorem 4 *Let $(\mathcal{K}, P, \{b\})$ be an \mathcal{ALCQI} learning instance. Then the following conditions are equivalent:*

1. $(\mathcal{K}, P, \{b\})$ has an \mathcal{ALCQI} solution;
2. there exists a model \mathcal{I} of \mathcal{K} such that for all $a \in P$, $\mathcal{A}, a \not\rightarrow^i \mathcal{I}, b$.

Now for the adaptation of Theorems 3 and 4 to DLs without inverse roles. Let $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ be an \mathcal{ALCQI} KB and $a \in \text{ind}(\mathcal{A})$. We use \mathcal{A}_a^\downarrow to denote the set of assertions in \mathcal{A} that use only individual names from the set $\text{reach}_{\mathcal{A}}(a)$ of individual names reachable from a in $G_{\mathcal{A}}$. Now $\mathcal{A}, a \rightarrow_r \mathcal{I}, b$ means that there is a homomorphism h from \mathcal{A}_a^\downarrow to \mathcal{I} with $h(a) = b$ such that the extended KB $\mathcal{K}_{\mathcal{I}, h}$, defined as

$$(\mathcal{T}, \mathcal{A} \cup \{C(c) \mid C \in \text{sub}(\mathcal{K}), c \in \text{reach}_{\mathcal{A}}(a), h(c) \in C^{\mathcal{I}}\}),$$

is satisfiable. We define $\mathcal{A}, a \rightarrow_r^i \mathcal{I}, b$ analogously, additionally requiring that h is locally injective for all role names r (but not necessarily for inverse roles).

Theorem 5 *Let $(\mathcal{K}, P, \{b\})$ be an \mathcal{ALC} (resp. \mathcal{ALCQ}) learning instance. Then the following conditions are equivalent:*

1. $(\mathcal{K}, P, \{b\})$ has an \mathcal{ALC} (resp. \mathcal{ALCQ}) solution;
2. there exists a model \mathcal{I} of \mathcal{K} such that for all $a \in P$, $\mathcal{A}, a \not\rightarrow_r \mathcal{I}, b$ (resp. $\mathcal{A}, a \not\rightarrow_r^i \mathcal{I}, b$).

The following example illustrates the reason for why the characterization in Theorem 5 requires $\mathcal{K}_{\mathcal{I}, h}$ to be satisfiable.

Example 2 Consider $\mathcal{A} = \{r(c, a), A(a), A(b)\}$, $P = \{a\}$, $N = \{b\}$, $\mathcal{T} = \{\top \sqsubseteq \forall r.B\}$, and let $\mathcal{K} = (\mathcal{T}, \mathcal{A})$. Then $\mathcal{K} \models B(a)$ and $\mathcal{K} \not\models B(b)$ and so B is an \mathcal{ALC} solution for (\mathcal{K}, P, N) . Observe that $h : a \mapsto b$ is a homomorphism from \mathcal{A}_a^\downarrow to any model \mathcal{I} of \mathcal{K} , but for any model \mathcal{I} of \mathcal{K} with $b \in (-B)^\mathcal{I}$ the extended KB $\mathcal{K}_{\mathcal{I},h}$ is not satisfiable.

We observe a slightly surprising consequence of the above characterizations. Concept separability is *TBox anti-monotone* in a DL \mathcal{L} if extending the TBox cannot result in a solution for a learning instance to become available, that is, for all \mathcal{L} learning instances (\mathcal{K}_1, P, N) and (\mathcal{K}_2, P, N) with $\mathcal{T}_1 \subseteq \mathcal{T}_2$ and $\mathcal{A}_1 = \mathcal{A}_2$, (\mathcal{K}_1, P, N) has an \mathcal{L} solution if (\mathcal{K}_2, P, N) has an \mathcal{L} solution.

Theorem 6 *Concept separability is TBox anti-monotone in \mathcal{ALCI} and \mathcal{ALCQI} , but not in \mathcal{EL} , \mathcal{ELI} , \mathcal{ALC} , and \mathcal{ALCQ} .*

Note that TBox anti-monotonicity in \mathcal{ALCI} and \mathcal{ALCQI} follows directly from Theorems 3 and 4. To show that concept separability is not TBox anti-monotone in \mathcal{ALC} and \mathcal{ALCQ} consider (\mathcal{K}, P, N) from Example 2 and let $\mathcal{K}' = (\emptyset, \mathcal{A})$. Then (\mathcal{K}', P, N) has no \mathcal{ALCQ} solution (and thus no \mathcal{ALC} solution), but (\mathcal{K}, P, N) has an \mathcal{ALC} solution. Example 1 (1) and (2) above shows that concept separability is not TBox anti-monotone in \mathcal{EL} and \mathcal{ELI} .

5 Complexity: Expressive DLs

We clarify the complexity of concept separability and concept definability in expressive DLs, which turn out to be NEXPTIME-complete in \mathcal{ALC} , \mathcal{ALCI} , and \mathcal{ALCQ} , and EXPTIME-complete in \mathcal{ALCQI} . In data complexity, they are Σ_2^p -complete in all four DLs. In both lower and upper bound proofs, we exploit a connection between separability and (appropriate versions of) UCQ evaluation on KBs that is suggested by the characterizations in the previous section.

Theorem 7 *\mathcal{ALCI} concept separability and \mathcal{ALCI} concept definability are NEXPTIME-complete.*

The upper bound in Theorem 7 is proved by a polynomial time reduction to the complement of rooted UCQ evaluation on \mathcal{ALCI} KBs and the lower bound is proved by a polynomial time reduction from the complement of rooted CQ evaluation on \mathcal{ALCI} KBs. Both problems are CONEXPTIME-complete [Lutz, 2008] and rather well known, so we refrain from giving definitions and only mention that a CQ is *rooted* if it is connected and has at least one answer variable, and that a UCQ is *rooted* if every CQ in it is.

For the first reduction, let $(\mathcal{K}, P, \{b\})$ be an \mathcal{ALCI} learning instance with $\mathcal{K} = (\mathcal{T}, \mathcal{A})$. For every $a \in P$, let q_a be the maximal connected component \mathcal{A}_a of \mathcal{A} that contains a , viewed as a connected CQ with only answer variable a . Further, let q be the UCQ $\bigvee_{a \in P} q_a$. Then $(\mathcal{K}, P, \{b\})$ has a solution iff $\mathcal{K} \not\models q(b)$. In fact, $\mathcal{K} \not\models q(b)$ iff there is a model \mathcal{I} of \mathcal{K} such that $\mathcal{I} \not\models q_a(b)$ for all $a \in P$, that is, $\mathcal{A}_a, a \not\rightarrow \mathcal{I}, b$. This is equivalent to $\mathcal{A}, a \not\rightarrow \mathcal{I}, b$ as required by Point 2 of Theorem 3 since for all other connected components \mathcal{A}' of \mathcal{A} , the identity is a homomorphism to \mathcal{I} .

For the second reduction, let $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ be an \mathcal{ALCI} KB, $q(x)$ a unary rooted CQ, and $a \in \text{ind}(\mathcal{A})$. Let \mathcal{A}_q be q viewed as an ABox where now x is an individual name, and let \mathcal{A}' be

the disjoint union of \mathcal{A} and \mathcal{A}_q . Let further $\mathcal{K}' = (\mathcal{T}, \mathcal{A}')$, $P = \{x\}$, and $N = \{a\}$. Using similar arguments as above, one can show that $\mathcal{K} \not\models q(a)$ iff (\mathcal{K}', P, N) has a solution. Note that this establishes hardness already for learning instances with a *single positive example* (and a single negative example). The reduction can be modified to show hardness of definability. In fact, the lower bound proof in [Lutz, 2008] applies already if we restrict ABoxes to be of the simple form $\{A_0(a_0)\}$ and can easily be modified so that the CQ $q(x)$ used is such that only the individual x can be an answer to q in \mathcal{A}_q (by introducing a fresh concept name X , adding the atom $X(x)$ to q and the assertion $X(a_0)$ to the ABox). Then we can set $P = \{x\}$ as before and N is the set of all other individuals in \mathcal{A}' .

In the case of \mathcal{ALCQI} , separability closely corresponds to a version of unary rooted (U)CQ evaluation that is based on locally injective homomorphisms from the CQ to models of the KB. We introduce this problem in the appendix and prove that, remarkably, it is only EXPTIME-complete. The intuitive reason is that the CONEXPTIME-lower bound for CQ evaluation on \mathcal{ALCI} KBs requires to ‘fold’ the CQ in exponentially many different ways, which can only be done with locally non-injective homomorphisms. In \mathcal{ALCQI} , concept separability is thus no harder than standard reasoning problems such as concept satisfiability w.r.t. a TBox—these are EXPTIME-complete for all expressive DLs considered in this paper [Baader *et al.*, 2017].

Theorem 8 *\mathcal{ALCQI} concept separability and \mathcal{ALCQI} concept definability are EXPTIME-complete.*

For the case without inverse roles, we also introduce corresponding versions of (U)CQ evaluation, based on the kinds of homomorphism used in Theorem 5. These turn out to be CONEXPTIME-complete. The lower bound is proved by a reduction of a suitable version of the tiling problem and crucially exploits the satisfiability condition used in the homomorphisms in Theorem 5. We work with a single positive example in the case of \mathcal{ALC} while in \mathcal{ALCQ} the number of positive examples is not bounded by a constant.

Theorem 9 *For $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCQ}\}$, \mathcal{L} concept separability and \mathcal{L} concept definability are NEXPTIME-complete.*

We close this section with clarifying data complexity.

Theorem 10 *For $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCQ}, \mathcal{ALCI}, \mathcal{ALCQI}\}$, \mathcal{L} concept separability and \mathcal{L} concept definability are Σ_2^p -complete in data complexity.*

The upper bounds are proved by a ‘guess-coguess and check’ style procedure. The lower bounds are proved by reduction of the problem to decide whether for a given undirected graph G and $k \geq 1$, there is a 2-coloring of G that does not generate a monochromatic k -clique [Rutenburg, 1986]. The proof requires two positive examples.

6 Complexity: Horn DLs

We study concept separability and concept definability in the case that the separation language is one of the Horn DLs \mathcal{EL} and \mathcal{ELI} . As the TBox language, we consider both Horn DLs and expressive DLs. It turns out that these problems tend to be

undecidable, notable exceptions being the cases that the TBox language is \mathcal{EL} or \mathcal{ELI} and the separation language is \mathcal{EL} .

We first refine the characterization from Theorem 2. For $\mathcal{L} \in \{\mathcal{EL}, \mathcal{ELI}\}$, a model \mathcal{U} of a knowledge base \mathcal{K} is \mathcal{L} universal if $\mathcal{K} \models C(a)$ iff $\mathcal{U} \models C(a)$, for all \mathcal{L} concepts C and $a \in \text{ind}(\mathcal{A})$.

Theorem 11 *Let $(\mathcal{K}, P, \{b\})$ be an \mathcal{ELI} learning instance, $\mathcal{L}_S \in \{\mathcal{EL}, \mathcal{ELI}\}$, and \mathcal{U} an \mathcal{L}_S universal model for \mathcal{K} of finite outdegree. Then the following are equivalent:*

1. $(\mathcal{K}, P, \{b\})$ has an \mathcal{L}_S solution,
2. $\prod_{a \in P} (\mathcal{U}, a) \not\leq_{\mathcal{L}_S} \mathcal{U}, b$.

The absence of disjunction in Horn DLs forces the product in Point 2, which captures the commonalities of the positive examples and achieves true generalization, avoiding overfitting. It is, however, also responsible for the high complexity and even undecidability of concept separability in Horn DLs. This is in fact rather surprising given that standard reasoning problems such as subsumption are only PTIME-complete in \mathcal{EL} and EXPTIME-complete in \mathcal{ELI} .

We start with the separation language \mathcal{EL} , considering both \mathcal{EL} and \mathcal{ELI} as the TBox language.

Theorem 12 *Let $\mathcal{L}_T \in \{\mathcal{EL}, \mathcal{ELI}\}$. Then $(\mathcal{L}_T, \mathcal{EL})$ concept separability and definability are EXPTIME-complete, both in combined complexity and in data complexity.*

For the lower bound, we adapt a proof used in [Harel *et al.*, 2002] to prove that the simulation problem between \mathcal{I} , d and \mathcal{J} , e is EXPTIME-hard when \mathcal{I} is a synchronized product of given interpretations and \mathcal{J} is an explicitly given interpretation. Remarkably, the lower bound holds already for empty TBoxes. The number of positive examples is not bounded by a constant. For the upper bound, we build on Point 2 of Theorem 11, constructing the product of (exponential size) \mathcal{EL} universal models of \mathcal{ELI} knowledge bases and then deciding \mathcal{EL} similarity in polynomial time. The upper bound easily extends to the case where the learned concept can use only symbols from a given signature.

Surprisingly, separability and definability even become undecidable when we extend \mathcal{EL} with inverse roles. This implies that no learning algorithm can be complete and terminating.

Theorem 13 *\mathcal{ELI} concept separability and definability are undecidable.*

The proof is by a rather subtle reduction of the tiling problem of rectangles of unbounded size, borrowing and extending some ideas from the proof given in [Botoeva *et al.*, 2019] that the CQ entailment problem between \mathcal{ALC} KBs is undecidable. It requires only two positive examples.

We remark that \mathcal{ELI} concept separability is closely related to the problem of learning conjunctive queries in the context of \mathcal{ELI} KBs ('query by example'). That problem was stated to be 2EXPTIME-complete in [Gutiérrez-Basulto *et al.*, 2018], but the proof of the upper bound turns out to be incorrect. In the appendix, we use a minor variation of our proof of Theorem 13 to show that the problem is actually undecidable.

Another related problem is the existence of most specific concepts (MSCs) and least common subsumers (LCSs) [Baader, 2003; Lutz *et al.*, 2010; Zarriß and Turhan,

2013]. In fact, if $(\mathcal{K}, P, \{b\})$ is a learning instance, MSCs M_a exist of all $a \in P$, and a LCS C of the M_a exists, then $(\mathcal{K}, P, \{b\})$ has a solution if $\mathcal{K} \not\models C(b)$. We conjecture that a variation of the proof of Theorem 13 can be used to show that in \mathcal{ELI} , the existence of the LCS is undecidable in the presence of a TBox.

We now consider the case where the TBox is formulated in an expressive DL and the separation language is a Horn DL. For example, such a setup is natural if avoiding overfitting is a concern. Also here, separability turns out to be undecidable.

Theorem 14 *Let $\mathcal{L}_T \in \{\mathcal{ALC}, \mathcal{ALCQ}, \mathcal{ALCI}, \mathcal{ALCQI}\}$ and $\mathcal{L}_S \in \{\mathcal{EL}, \mathcal{ELI}\}$. Then $(\mathcal{L}_T, \mathcal{L}_S)$ concept separability and $(\mathcal{L}_T, \mathcal{L}_S)$ concept definability are undecidable.*

The proof is by reduction from the already mentioned CQ entailment problem between \mathcal{ALC} KBs which undecidable already for (directed or undirected) tree-shaped CQs.

7 First Observations on Strong Separability

There is an alternative notion of concept learning that has been proposed in the literature [Lehmann *et al.*, 2014; Badae and Nienhuys-Cheng, 2000; Fanizzi *et al.*, 2018]. An \mathcal{L} concept C is a *strong \mathcal{L} solution* to a learning instance (\mathcal{K}, P, N) if $\mathcal{K} \models C(a)$ for all $a \in P$ and $\mathcal{K} \models \neg C(a)$ for all $a \in N$. We then define *strong \mathcal{L} concept separability* in the expected way. While we leave a thorough investigation for future work, we observe the following characterization.

Theorem 15 *For $\mathcal{L}_S \in \{\mathcal{ALC}, \mathcal{ALCI}, \mathcal{ALCQ}, \mathcal{ALCQI}\}$, an \mathcal{ALCQI} learning instance (\mathcal{K}, P, N) has a strong \mathcal{L} solution iff for all models \mathcal{I} and \mathcal{J} of \mathcal{K} , all $a \in P$ and all $b \in N$, $\mathcal{I}, a \not\leq_{\mathcal{L}_S} \mathcal{J}, b$.*

Based on this result, one can show that for \mathcal{ALC} and \mathcal{ALCI} there is a tight link between strong concept separability and KB satisfiability (rather than to UCQ-evaluation on KBs) which can be used to prove the following.

Theorem 16 *For $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCI}\}$, strong \mathcal{L} concept separability is EXPTIME-complete in combined complexity and CONP-complete in data complexity.*

Note the drop in complexity compared to the non-strong version of concept separability, from NEXPTIME to EXPTIME and from Σ_2^P to CONP. For \mathcal{ALCQ} and \mathcal{ALCQI} , the complexity of strong \mathcal{L} concept separability remains open.

8 Discussion

This paper provides characterizations and complexity results for concept separability in many important DLs. It would be interesting to further add other popular DL features such as role hierarchies, transitive roles, and nominals, and it might also be relevant to consider cases of $(\mathcal{L}_T, \mathcal{L}_S)$ concept separability in which \mathcal{L}_T is less expressive than \mathcal{L}_S . We also plan to investigate strong separability and in how the characterizations given in this paper can be used or extended to analyse the version space beyond emptiness and whether they can also be used to analyse or craft refinement operators.

Acknowledgements

Lutz was supported by DFG CRC 1320 Ease.

References

- [Baader *et al.*, 2017] Franz Baader, Ian Horrocks, Carsten Lutz, and Ulrike Sattler. *An Introduction to Description Logic*. Cambridge University Press, 2017.
- [Baader, 2003] Franz Baader. Least common subsumers and most specific concepts in a description logic with existential restrictions and terminological cycles. In *Proc. of IJCAI*, pages 319–324, 2003.
- [Badea and Nienhuys-Cheng, 2000] Liviu Badea and Shan-Hwei Nienhuys-Cheng. A refinement operator for description logics. In *Proc. of ILP*, pages 40–59, 2000.
- [Bienvenu *et al.*, 2014] Meghyn Bienvenu, Balder ten Cate, Carsten Lutz, and Frank Wolter. Ontology-based data access: A study through disjunctive datalog, CSP, and MM-SNP. *ACM Trans. Database Syst.*, 39(4):33:1–33:44, 2014.
- [Botoeva *et al.*, 2019] Elena Botoeva, Roman Kontchakov, Carsten Lutz, Vladislav Ryzhikov, Frank Wolter, and Michael Zakharyashev. Query inseparability for ALC ontologies. *Artif. Intell.*, 272:1–51, 2019.
- [Bühmann *et al.*, 2016] Lorenz Bühmann, Jens Lehmann, and Patrick Westphal. DL-learner - A framework for inductive learning on the semantic web. *J. Web Sem.*, 39:15–24, 2016.
- [Bühmann *et al.*, 2018] Lorenz Bühmann, Jens Lehmann, Patrick Westphal, and Simon Bin. DL-learner structured machine learning on semantic web data. In *Proc. of WWW*, pages 467–471, 2018.
- [Calvanese *et al.*, 2013] Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. Data complexity of query answering in description logics. *Artif. Intell.*, 195:335–360, 2013.
- [Fanizzi *et al.*, 2018] Nicola Fanizzi, Giuseppe Rizzo, Claudia d’Amato, and Floriana Esposito. DLFOIL: Class expression learning revisited. In *Proc. of EKAW*, pages 98–113, 2018.
- [Goranko and Otto, 2007] Valentin Goranko and Martin Otto. Model theory of modal logic. In *Handbook of Modal Logic*, pages 249–329. Elsevier, 2007.
- [Gutiérrez-Basulto *et al.*, 2018] Víctor Gutiérrez-Basulto, Jean Christoph Jung, and Leif Sabellek. Reverse engineering queries in ontology-enriched systems: The case of expressive horn description logic ontologies. In *Proc. of IJCAI-ECAI*, 2018.
- [Harel *et al.*, 2002] David Harel, Orna Kupferman, and Moshe Y. Vardi. On the complexity of verifying concurrent transition systems. *Inf. Comput.*, 173(2):143–161, 2002.
- [Hirsh *et al.*, 2004] Haym Hirsh, Nina Mishra, and Leonard Pitt. Version spaces and the consistency problem. *Artif. Intell.*, 156(2):115–138, 2004.
- [Iannone *et al.*, 2007] Luigi Iannone, Ignazio Palmisano, and Nicola Fanizzi. An algorithm based on counterfactuals for concept learning in the semantic web. *Appl. Intell.*, 26(2):139–159, 2007.
- [Kietz, 1993] Jörg-Uwe Kietz. Some lower bounds for the computational complexity of inductive logic programming. In *Proc. of ECML*, pages 115–123, 1993.
- [Lehmann and Haase, 2009] Jens Lehmann and Christoph Haase. Ideal downward refinement in the \mathcal{EL} description logic. In *Proc. of ILP*, pages 73–87, 2009.
- [Lehmann and Hitzler, 2010] Jens Lehmann and Pascal Hitzler. Concept learning in description logics using refinement operators. *Machine Learning*, 78:203–250, 2010.
- [Lehmann and Völker, 2014] Jens Lehmann and Johanna Völker. *Perspectives on Ontology Learning*, volume 18 of *Studies on the Semantic Web*. IOS Press, 2014.
- [Lehmann *et al.*, 2014] Jens Lehmann, Nicola Fanizzi, Lorenz Bühmann, and Claudia d’Amato. Concept learning. In *Perspectives on Ontology Learning*, pages 71–91. AKA / IOS Press, 2014.
- [Lisi, 2012] Francesca A. Lisi. A formal characterization of concept learning in description logics. In *Proc. of DL*, 2012.
- [Lutz *et al.*, 2010] Carsten Lutz, Robert Piro, and Frank Wolter. Enriching \mathcal{EL} -concepts with greatest fixpoints. In *Proc. of ECAI*, pages 41–46, 2010.
- [Lutz *et al.*, 2011] Carsten Lutz, Robert Piro, and Frank Wolter. Description logic TBoxes: Model-theoretic characterizations and rewritability. In *Proc. of IJCAI*, 2011.
- [Lutz, 2008] Carsten Lutz. The complexity of conjunctive query answering in expressive description logics. In *Proc. of IJCAR*, pages 179–193, 2008.
- [Pitt and Valiant, 1988] Leonard Pitt and Leslie G. Valiant. Computational limitations on learning from examples. *J. ACM*, 35(4):965–984, 1988.
- [Rutenburg, 1986] Vladislav Rutenburg. Complexity of generalized graph coloring. In *Proc. of MFCS*, pages 573–581, 1986.
- [Sarker *et al.*, 2017] Md. Kamruzzaman Sarker, Ning Xie, Derek Doran, Michael Raymer, and Pascal Hitzler. Explaining trained neural networks with semantic web technologies: First steps. In *Proc. of NeSy*, 2017.
- [Straccia and Mucci, 2015] Umberto Straccia and Matteo Mucci. pFOIL-DL: Learning (fuzzy) \mathcal{EL} concept descriptions from crisp OWL data using a probabilistic ensemble estimation. In *Proc. of SAC*, pages 345–352. ACM, 2015.
- [Tran *et al.*, 2014] Thanh-Luong Tran, Quang-Thuy Ha, Thi-Lan-Giao Hoang, Linh Anh Nguyen, and Hung Son Nguyen. Bisimulation-based concept learning in description logics. *Fundam. Inform.*, 133(2-3):287–303, 2014.
- [Zarriëß and Turhan, 2013] Benjamin Zarriëß and Anni-Yasmin Turhan. Most specific generalizations w.r.t. general \mathcal{EL} -tboxes. In *Proc. of IJCAI*, pages 1191–1197, 2013.