

SAT-Based Bounded Fitting for the Description Logic \mathcal{ALC}

Maurice Funk¹[0000-0003-1823-9370], Jean Christoph Jung²[0000-0002-4159-2255],
and Tom Voellmer²[0009-0006-3149-7066]

¹ Leipzig University and ScaDS.AI, mfunk@informatik.uni-leipzig.de

² TU Dortmund University, {jean.jung, tom.voellmer}@tu-dortmund.de

Abstract. Bounded fitting is a general paradigm for learning logical formulas from positive and negative data examples, that has received considerable interest recently. We investigate bounded fitting for the description logic \mathcal{ALC} and its syntactic fragments. We show that the underlying size-restricted fitting problem is NP-complete for all studied fragments, even in the special case of a single positive and a single negative example. By design, bounded fitting comes with probabilistic guarantees in Valiant’s PAC learning framework. In contrast, we show that other classes of algorithms for learning \mathcal{ALC} concepts do not provide such guarantees. Finally, we present an implementation of bounded fitting in \mathcal{ALC} and its fragments based on a SAT solver. We discuss optimizations and compare our implementation to other concept learning tools.

Keywords: Class Expression Learning · Description Logic · SAT solving

1 Introduction

Learning *class expressions* from given data examples is an important task when working with large knowledge bases [25, 26]. For the purpose of this paper, an *example* is a pair (\mathcal{I}, a) where \mathcal{I} is a data instance, e.g., a database or a knowledge graph, and a is some element in the instance. Moreover, a class expression C *fits* a set P of *positive examples* and a set N of *negative examples* if $\mathcal{I} \models C(a)$ for all $(\mathcal{I}, a) \in P$ and $\mathcal{I} \not\models C(a)$ for all $(\mathcal{I}, a) \in N$. We mention three applications. First, the fitting class expression may be used as an *explanation* of the separation between good and bad scenarios, described by P and N , respectively. For example, P and N could be data describing users who visited (resp., did not visit) a certain web page, and a fitting C would explain the users’ behavior based on their data. Second, under the classical *query-by-example* paradigm [28, 40], a human user may *reverse-engineer* a class expression by manually selecting elements they want to have returned (P) or not returned (N), and a system produces an expression satisfying the demands. Here, the intended use of the fitting class expression is as a query. Finally, an ontology engineer may seek a definition of some class A occurring in the data instance, so they may request a class expression C separating all elements P satisfying A from those that do not (N), and add $A \equiv C$ to the ontology.

Input: Positive examples P , negative examples N

```

1 for  $k := 1, 2, \dots$  do
2   | if there is  $\varphi \in \mathcal{L}$  of size  $k$  that fits  $P, N$  then
3   |   | return  $\varphi$ 

```

Algorithm 1: Bounded Fitting for abstract logic \mathcal{L} .

In this paper, we study the problem of learning class expressions formulated in the description logic (DL) \mathcal{ALC} . \mathcal{ALC} can be viewed as the conceptual basis of the very expressive description logic $\mathit{SROIQ}(\mathcal{D})$ that underlies the web ontology language OWL 2 DL [1, 17] in that \mathcal{ALC} provides the central logical constructors that are also available in $\mathit{SROIQ}(\mathcal{D})$. The importance of finding fitting class expressions / description logic concepts has resulted in both foundational work [14, 15, 27] and implemented systems. While most systems are based on heuristic search and refinement operators [6, 16, 18, 34, 37] or, more recently, also on neural techniques [13, 23], we approach the problem via bounded fitting. *Bounded fitting* is a general paradigm for fitting logical formulas to positive and negative examples that has been investigated recently for a range of logics like linear temporal logic LTL [7, 30], computational tree logic CTL [32], and the description logic \mathcal{EL} [9]. Algorithm 1 provides an abstract description of bounded fitting for a given logic \mathcal{L} . It should be clear that, if any fitting formula exists, bounded fitting always returns a fitting formula of minimal size, which is often a desirable property. From a practical perspective, human users typically prefer shorter, that is, simpler formulas in the applications described above. For instance, when class expressions are used as explanations, a shorter class expression might provide a more transparent explanation. From a theoretical perspective, this property makes bounded fitting an *Occam algorithm*, which implies that it comes with probabilistic generalization guarantees in Valiant’s probably approximately correct (PAC) learning framework [4, 38]. Intuitively, this means that bounded fitting needs only few examples to be able to generalize to unseen examples. A further advantage of bounded fitting is that it is *complete*: whenever there is a fitting concept bounded fitting will find one.

The basic DL \mathcal{ALC} provides the logical constructors conjunction \sqcap , disjunction \sqcup , negation \neg , existential restriction $\exists r$, and universal restriction $\forall r$ to build complex concepts. Motivated by the fact that, depending on the application, one may not need all concept constructors, fragments $\mathcal{L}(O)$ of \mathcal{ALC} have been studied which allow only a subset $O \subseteq \{\sqcap, \sqcup, \neg, \exists, \forall\}$ of the available constructors. For instance, the mentioned DL \mathcal{EL} is defined by $\{\sqcap, \exists\}$ and the DL \mathcal{FL}_0 is defined by $\{\sqcap, \forall\}$. Bounded fitting has been studied recently for \mathcal{EL} , and we extend this study here to all other syntactical fragments.

Our main contributions are as follows. First, we study the *size-restricted fitting problem*: given positive examples P , negative examples N , and a size bound k , determine whether there is a concept of size at most k that fits P and N . Clearly, this is precisely the problem to be solved in Line 2 of Algorithm 1. Our main result is that size-restricted fitting for $\mathcal{L}(O)$ is NP-complete for all O that include at least \exists or \forall . While the NP upper bound is straightforward, the

lower bound is more technical and rather strong. First, surprisingly, it already applies for inputs consisting of one positive and one negative example, and thus strengthens known hardness results for \mathcal{EL} [10,11]. Second, it uses only a fixed set of symbols in contrast to related hardness results for temporal logics [5,29].

Then, motivated by the ability of bounded fitting to generalize well from few examples, we discuss the generalization abilities of fitting algorithms for languages $\mathcal{L}(O)$ in Valiant's PAC learning framework. We start with observing that under reasonable complexity theoretic assumptions, *no* $\mathcal{L}(O)$ admits an *efficient* PAC learning algorithm. We then analyze the generalization ability of fitting algorithms that have favorable properties from a logical perspective in that they return fitting formulas that are most specific, most general, or of minimal quantifier depth among all fitting formulas. We show that, with one exception, all such algorithms are not *sample-efficient*, and hence do not generalize well. Similar results have been obtained for \mathcal{EL} concepts [9] and also for the related query language of conjunctive queries [8].

Finally, we provide an implementation of bounded fitting for \mathcal{ALC} and its fragments, that relies on a SAT solver to solve size-restricted fitting in Line 2 of bounded fitting. We describe our implementation and present several improvements in the encoding which lead to significant speed-ups. We also present an *approximation scheme*, which is an adaptation of the basic bounded fitting algorithm depicted in Algorithm 1 to be able to provide approximate solutions in case no perfect fitting was found. We then compare our implementation to other concept learning tools from the literature, both in terms of achieved accuracy, size of the returned concepts, and generalization ability.

2 Preliminaries

We introduce syntax and semantics of the description logic \mathcal{ALC} [3]. Let \mathbb{N}_C and \mathbb{N}_R be mutually disjoint and countably infinite sets of *concept names* and *role names*, respectively. An \mathcal{ALC} *concept* C is defined according to the syntax rule

$$C, D ::= \perp \mid \top \mid A \mid \neg C \mid C \sqcap D \mid C \sqcup D \mid \forall r.C \mid \exists r.C$$

where A ranges over concept names and r over role names. For any set $O \subseteq \{\neg, \sqcap, \sqcup, \forall, \exists\} = O_{\text{all}}$, we define $\mathcal{L}(O)$ as the set of all \mathcal{ALC} concepts built only from connectives and quantifiers from O . We associate with each such set O a *dual* set \bar{O} that contains precisely the symbols dual to those in O where \neg is dual to itself, \sqcap is dual to \sqcup , and \exists is dual to \forall . The size $\|C\|$ of an \mathcal{ALC} concept C is defined as the number of nodes in its syntax tree, see below for an example. A *signature* Σ is a finite set of concept and role names. The *signature of an \mathcal{ALC} concept* C is the set of concept and role names that occur in C .

The semantics of \mathcal{ALC} concepts is given as usual in terms of interpretations. An *interpretation* is a pair $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ where $\Delta^{\mathcal{I}}$ is a non-empty set of elements, called *domain*, and $\cdot^{\mathcal{I}}$ is an interpretation function that assigns a set $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$ to every $A \in \mathbb{N}_C$ and a binary relation $r^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$ to every $r \in \mathbb{N}_R$. The

extension $C^{\mathcal{I}}$ of complex \mathcal{ALC} concepts C is defined inductively by taking $\perp^{\mathcal{I}} = \emptyset$, $\top^{\mathcal{I}} = \Delta^{\mathcal{I}}$, $(\neg C)^{\mathcal{I}} = \Delta^{\mathcal{I}} \setminus C^{\mathcal{I}}$, $(C \sqcap D)^{\mathcal{I}} = C^{\mathcal{I}} \cap D^{\mathcal{I}}$, $(C \sqcup D)^{\mathcal{I}} = C^{\mathcal{I}} \cup D^{\mathcal{I}}$, and

$$(\exists r.C)^{\mathcal{I}} = \{d \in \Delta^{\mathcal{I}} \mid \text{there is } (d, e) \in r^{\mathcal{I}} \text{ with } e \in C^{\mathcal{I}}\}$$

and $(\forall r.C)^{\mathcal{I}} = (\neg \exists r. \neg C)^{\mathcal{I}}$.

We call an interpretation \mathcal{I} *finite*, if both $\Delta^{\mathcal{I}}$ and the set of all concept and role names that have a nonempty extension in \mathcal{I} are finite. A finite interpretation \mathcal{I} corresponds to a finite set of facts

$$\{A(a) \mid A \in \mathbf{N}_C, a \in A^{\mathcal{I}}\} \cup \{r(a, b) \mid r \in \mathbf{N}_R, (a, b) \in r^{\mathcal{I}}\},$$

and can hence be viewed as a database. Conversely, every database or finite set of facts corresponds to a finite interpretation. We define an example as a pair (\mathcal{I}, a) consisting of a finite interpretation \mathcal{I} and some domain element $a \in \Delta^{\mathcal{I}}$. Let P and N be sets of examples which we will refer to as *positive* and *negative* examples. We say that an \mathcal{ALC} concept C fits P, N if $a \in C^{\mathcal{I}}$ for each $(\mathcal{I}, a) \in P$ and $b \notin C^{\mathcal{J}}$ for each $(\mathcal{J}, b) \in N$. In the special case of $P = \{(\mathcal{I}, a_0)\}$ and $N = \{(\mathcal{J}, b_0)\}$ we just write that C fits $(\mathcal{I}, a_0), (\mathcal{J}, b_0)$.

The *size* of such an example (\mathcal{I}, a) , denoted $\|(\mathcal{I}, a)\|$, is $n + 1$ where n is the cardinality of \mathcal{I} , viewed as set of facts. As usual, we view interpretations as node and edge labeled graphs and use graph terminology to speak about them.

We investigate the following decision problem.

Problem: SIZE-RESTRICTEDFITTING FOR \mathcal{L}

Given: Sets P, N of examples, integer k

Question: Is there an \mathcal{L} concept C that fits P, N and has size $\|C\| \leq k$?

We study size-restricted fitting for syntactic fragments $\mathcal{L}(O)$ of \mathcal{ALC} , assuming that k is given in unary, which is a natural choice in the context of bounded fitting. If we drop the size restriction k from the input, we obtain the *fitting problem for \mathcal{L}* .

Example 1. As an example for size-restricted fitting, consider the two positive examples $P = \{(\mathcal{I}, a_1), (\mathcal{I}, a_2)\}$ and one negative example $N = \{(\mathcal{J}, b)\}$ depicted in Figure 1. Then $P, N, k = 4$ is a yes-instance of size-restricted fitting for \mathcal{ALC} with witnessing fitting concept $C = \forall r.(A \sqcup B)$ with $\|\forall r.(A \sqcup B)\| = 4$, see the syntax tree depicted in the right of Figure 1. Indeed, we have $a_1, a_2 \in C^{\mathcal{I}}$ but $b \notin C^{\mathcal{J}}$. It can also be verified, that it is the minimal k for which there exists a fitting concept, and that there is no fitting $\mathcal{L}(\{\exists, \sqcap\})$ concept at all, hence P, N is a no-instance for the fitting problem for $\mathcal{L}(\{\exists, \sqcap\})$. \dashv

To transfer results between dual fragments of \mathcal{ALC} , we rely on an auxiliary lemma that relates fitting in $\mathcal{L}(O)$ to fitting in $\mathcal{L}(\bar{O})$. To formulate the lemma, we define a *duality* mapping $\bar{\cdot}$ for \mathcal{ALC} concepts, interpretations and sets of examples. For an \mathcal{ALC} concept C , the dual concept \bar{C} is obtained from C by replacing \top with \perp and vice versa, replacing \sqcap with \sqcup and vice versa, and replacing \exists with \forall and vice versa. Moreover, for an interpretation \mathcal{I} and a signature Σ , the




Fig. 1: Example for size-restricted fitting

interpretation $\bar{\mathcal{I}}_\Sigma$ is obtained by complementing the interpretation of the concept names in Σ , that is, $\Delta^{\bar{\mathcal{I}}_\Sigma} = \Delta^{\mathcal{I}}$,

$$A^{\bar{\mathcal{I}}_\Sigma} = \begin{cases} \Delta^{\mathcal{I}} \setminus A^{\mathcal{I}} & \text{if } A \in \mathbf{N}_C \cap \Sigma, \\ A^{\mathcal{I}} & \text{if } A \in \mathbf{N}_C \setminus \Sigma, \end{cases}$$

and keeping the interpretation of the role names, that is, $r^{\bar{\mathcal{I}}_\Sigma} = r^{\mathcal{I}}$, for all $r \in \mathbf{N}_R$. One can then show inductively that the following duality is satisfied for all signatures Σ , \mathcal{ALC} concepts C with signature contained in Σ , interpretations \mathcal{I} , and $a \in \Delta^{\mathcal{I}}$:

$$a \in C^{\mathcal{I}} \quad \text{if and only if} \quad a \notin \bar{C}^{\bar{\mathcal{I}}_\Sigma}. \quad (1)$$

For a set of examples Ex and signature Σ , we define the set $\bar{Ex}_\Sigma = \{(\bar{\mathcal{I}}_\Sigma, a) \mid (\mathcal{I}, a) \in Ex\}$. Duality (1) is used to relate fitting in dual languages as follows:

Lemma 1. *For any signature Σ , all sets of examples P and N and \mathcal{ALC} concepts C with signature contained in Σ , C fits P, N if and only if \bar{C} fits $\bar{N}_\Sigma, \bar{P}_\Sigma$.*

Since for given sets of examples P, N there is a fitting \mathcal{L} concept of size k if and only if there is a fitting \mathcal{L} concept of size k which uses only concept and role names that occur in P, N , Lemma 1 implies the following:

Corollary 1. *For every $O \subseteq O_{all}$, there is a polynomial time reduction of size-restricted fitting for $\mathcal{L}(O)$ to size-restricted fitting for $\mathcal{L}(\bar{O})$.*

3 Complexity of Size-Restricted Fitting

We determine the complexity of size-restricted fitting as the decision problem underlying bounded fitting in all syntactic fragments $\mathcal{L}(O)$ where $O \subseteq O_{all}$ and $O \cap \{\forall, \exists\} \neq \emptyset$. Our main result states that the problem is NP-complete, for any such O . While the proof of the upper bound is straightforward, the NP-hardness proof is rather involved. Notably, it holds already for single positive and negative examples, which came as a surprise to us.

Theorem 1. *For every $O \subseteq O_{all}$ such that O contains at least one of \exists and \forall , size-restricted fitting for $\mathcal{L}(O)$ is NP-complete.*

Containment in NP follows from a standard guess-and-check argument. Given P, N and k , the algorithm non-deterministically picks an $\mathcal{L}(O)$ concept C of size at most k and verifies (in polynomial time) that $a \in C^{\mathcal{I}}$ for all $(\mathcal{I}, a) \in P$ and $b \notin C^{\mathcal{J}}$ for all $(\mathcal{J}, b) \in N$. Showing hardness is more involved and proceeds in two steps. First, with the following proposition, we show NP-hardness for any O with $\exists \in O$. Then, Corollary 1 implies NP-hardness for any O with $\forall \in O$.

Proposition 1. *Size-restricted fitting for $\mathcal{L}(O)$ is NP-hard for any $O \subseteq O_{all}$ with $\exists \in O$. This already holds if only a single positive and a single negative example are allowed, and over a signature consisting of two role names and one concept name.*

Proof. We give a reduction of *Hitting Set* which is the problem to decide, given a collection of sets $S = \{S_1, \dots, S_m\}$ and a size bound $k \in \mathbb{N}$, whether there exists a set H such that $|H| \leq k$ and $H \cap S_j \neq \emptyset$, for each $j \in \{1, \dots, m\}$. The set H is called a hitting set for S .

Let (S, k) be an instance of *Hitting Set* with $S = \{S_1, \dots, S_m\}$ a collection of sets $S_i \subseteq \mathbb{N}$ and $k \in \mathbb{N}$. Assume without loss of generality that $\bigcup_{j=1}^m S_j = \{1, \dots, n\}$ for some $n \in \mathbb{N}$. We construct a positive example (\mathcal{I}, a) and a negative example (\mathcal{J}, b) in which we interpret only two role names r, s and a concept name A non-empty, and such that, for $P = \{(\mathcal{I}, a)\}$, $N = \{(\mathcal{J}, b)\}$, and $k' = k + n + 2$,

$$S \text{ has hitting set of size } k \iff P, N \text{ admit an } \mathcal{L}(O) \text{ fitting of size } k'. \quad (2)$$

We provide an informal description of \mathcal{I}, \mathcal{J} ; more details can be found in the appendix. The main components of both \mathcal{I} and \mathcal{J} are r -paths of length n , which end in an element satisfying A . To encode a subset $S' \subseteq \{1, \dots, n\}$, such a path is extended with additional “detours” along role s . In more detail, suppose such a path consists of elements b_0, \dots, b_n . We encode a subset $S' \subseteq \{1, \dots, n\}$ as follows:

- For all $i \in \{1, \dots, n\} \setminus S'$, there is an s -path of length 2 from b_i to b_{i+1} .
- For every b_i and role $t \in \{r, s\}$: if b_i does not have a t -successor, then make c a t -successor of b_i where c is a *sink*, which has both an r and an s self loop.

Now, \mathcal{J} consists of one path for each $S_j \in S$ and an additional root node b which connects via r to the initial nodes of all the paths. The interpretation \mathcal{I} is an extension of \mathcal{J} with an additional component that encodes the empty set, and has root node a which connects via r to the initial nodes of all paths. Figure 2 depicts the constructed interpretations for the instance $(S, 2)$ for $S = \{\{1, 3\}, \{2, 4\}\}$. We omit the sink element c in the figure for the sake of readability. For this input, $H = \{1, 2\}$ is a hitting set of size 2.

We verify Equivalence (2) by showing that the following are equivalent:

- (i) S has a hitting set of size at most k ;
- (ii) there is an $\mathcal{L}(\{\exists\})$ concept fitting P, N of size at most k' ;
- (iii) there is an \mathcal{ALC} concept fitting P, N of size at most k' .




Fig. 2: Example of the interpretations \mathcal{I} and \mathcal{J} used in the reduction.

We start with (i) \Rightarrow (ii). Let H be a hitting set for $S = \{S_1, \dots, S_m\}$ with $|H| = k$. We inductively define concepts C_i , for $i = 0, \dots, n$, by setting $C_0 = A$ and

$$C_i = \begin{cases} \exists r.C_{i-1} & \text{if } n - i + 1 \notin H \\ \exists s.\exists s.C_{i-1} & \text{otherwise,} \end{cases} \quad (3)$$

for $1 \leq i \leq n$. Thus, $\|C_n\| = n + |H| + 1$ and therefore $\|\exists r.C_n\| = n + k + 2 = k'$. We claim that $D = \exists r.C_n$ fits P, N . In the example in Figure 2, the concept constructed for the hitting set $H = \{1, 2\}$ is $D = \exists r.\exists s.\exists s.\exists s.\exists s.\exists r.\exists r.A$, and it is easily verified that it fits. In general, any concept that can be obtained this way fits the positive example (\mathcal{I}, a) . Moreover, a concept D constructed from a hitting set H fits the negative example (\mathcal{J}, b) since the construction forces the path described by D to “leave” the paths encoding the sets S_j in (\mathcal{J}, b) . This means that D cannot be satisfied in b , since the sink c does not satisfy A .

The direction (ii) \Rightarrow (iii) is immediate. For (iii) \Rightarrow (i), we take advantage of the structure of (\mathcal{I}, a) and (\mathcal{J}, b) to show that it is without loss of generality to assume that the smallest fitting is from $\mathcal{L}(\{\exists\})$ and has the shape from Equation (3). The main properties we exploit are that, except at the roots a, b , r, s are interpreted as total functions (to replace $\forall r/s$ by $\exists r/s$) and that there are single positive and negative examples (to remove disjunction and conjunction, respectively). \square

4 Generalization

In this section, we investigate the generalization ability of fitting algorithms with certain properties using the well-known PAC learning framework [38], see also [2] for an introduction. Here, with a *fitting algorithm for \mathcal{L}* , we mean an algorithm \mathcal{A} that receives as input P, N and returns an \mathcal{L} concept C that fits P, N if such a concept exists. To define PAC learning, we need some additional notation.

A *probability distribution over examples* is a function \mathbb{P} that maps examples (\mathcal{I}, a) to a non-negative value $\mathbb{P}(\mathcal{I}, a)$ in a way such that $\sum_{(\mathcal{I}, a)} \mathbb{P}(\mathcal{I}, a) = 1$. A *signature* is a finite set Σ of concept and role names. We say that \mathbb{P} is *over* Σ if in examples (\mathcal{I}, a) with $\mathbb{P}(\mathcal{I}, a) > 0$, the \mathcal{I} contains only facts $A(a)$ and $r(a, b)$ with $A \in \Sigma$ and $r \in \Sigma$, respectively. A *sample* of \mathbb{P} is a set Ex of examples drawn according to \mathbb{P} . *Labeling Ex according to a (target) concept C_T* means to split Ex into P, N by taking $P = \{(\mathcal{I}, a) \in Ex \mid a \in C_T^{\mathcal{I}}\}$ and $N = Ex \setminus P$. The *error* of a (hypothesis) concept C_H compared to C_T is defined as the probability that C_H and C_T label a random example differently, that is:

$$\text{error}_{\mathbb{P}}(C_T, C_H) = \mathbb{P}(\{(\mathcal{I}, a) \mid a \in C_T^{\mathcal{I}} \Delta C_H^{\mathcal{I}}\})$$

where Δ denotes symmetric difference. We denote with $(0, 1)$ the open interval between 0 and 1.

Definition 1. *Let $O \subseteq O_{all}$. A fitting algorithm \mathcal{A} for $\mathcal{L}(O)$ is a PAC learning algorithm if there is a function $m: \mathbb{R}^2 \times \mathbb{N}^3 \rightarrow \mathbb{N}$, such that for all $\varepsilon, \delta \in (0, 1)$, all signatures Σ , all $s, n \in \mathbb{N}$, all probability distributions \mathbb{P} over examples (\mathcal{I}, a) over Σ with $\|(\mathcal{I}, a)\| \leq s$, and all target concepts $C_T \in \mathcal{L}(O)$ over Σ with $\|C_T\| \leq n$, algorithm \mathcal{A} has the following property. If \mathcal{A} receives $m(\frac{1}{\delta}, \frac{1}{\varepsilon}, |\Sigma|, s, n)$ examples P, N sampled from \mathbb{P} and labeled according to C_T , it returns with high probability at least $1 - \delta$ a concept $C_H \in \mathcal{L}(O)$ with small $\text{error}_{\mathbb{P}}(C_T, C_H) \leq \varepsilon$.*

A PAC learning algorithm is said to be *sample-efficient* if m is a polynomial, and *efficient* if it is sample-efficient and runs in polynomial time.

Unfortunately, efficient PAC learnability is known to be elusive, and we confirm this to be the case here as well, under reasonable assumptions.

Theorem 2. *Let $O \subseteq O_{all}$. If there is an efficient PAC learning algorithm for $\mathcal{L}(O)$, then:*

1. $\text{NP} = \text{RP}$, if O contains at least one of \exists/\forall and $\{\sqcap, \sqcup\} \not\subseteq O$;
2. *RSA encryption is polynomial time invertible, if $\{\sqcap, \sqcup\} \subseteq O$.*

Recall that RP is the class of problems for which probabilistic polynomial-time Turing machines exist that reject no-instances, and accept yes instances with probability $\frac{1}{2}$. Both $\text{NP} = \text{RP}$ and polynomial time RSA decryption (without private key) are unlikely. Point 1 was known for $\mathcal{EL} = \mathcal{L}(\{\exists, \sqcap\})$ [11, 22] and we show it to be the case for all other mentioned fragments. The proof relies on a well-known result due to Pitt and Valiant which roughly states that efficient PAC learnability implies (under suitable conditions) that the corresponding (not size restricted) fitting problem is in RP [31]. To show Theorem 2, we thus consider the relevant fitting problems. In the appendix, we show that the fitting problem for $\mathcal{L}(\{\forall\})$, $\mathcal{L}(\{\forall, \sqcap\})$, and $\mathcal{L}(\{\forall, \exists, \sqcap\})$ is NP -hard via reduction from SAT. Applying Lemma 1 then yields the same lower bounds for all other fragments from Point 1. The same strategy fails for O as in Point 2 as the fitting problem for such O is solvable in polynomial time. Instead, we observe that $\mathcal{L}(O)$ contains

all monotone Boolean formulas and exploit that these are not efficiently PAC learnable under cryptographic assumptions [21].

Given these negative results on *efficient* PAC learning, we will analyze concrete fitting algorithms regarding their sample-efficiency. More precisely, we study algorithms that behave well from a logical perspective in that they return fitting concepts that are most specific, most general, or of minimal quantifier depth among all fitting concepts. We denote with $C \sqsubseteq D$ the fact that $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$, for all interpretations \mathcal{I} , and let $C \equiv D$ abbreviate $C \sqsubseteq D$ and $D \sqsubseteq C$. Recall that a concept C is a *most specific fitting* for P, N if it fits P, N and there is no concept D with $D \sqsubseteq C$ and $D \neq C$ that fits. Dually, a concept C is a *most general fitting* concept for P, N if it fits P, N and there exists no D with $C \sqsubseteq D$ and $C \neq D$ that fits as well. Finally, the *quantifier depth* of a concept is the nesting depth of \exists/\forall . Unfortunately, in most cases, logically well-behaved algorithms are not sample-efficient. The following theorem summarizes the situation.

Theorem 3. *Let $O \subseteq O_{all}$ be any set containing at least one of \exists/\forall and at least one of \sqcap/\sqcup , and let \mathcal{A} be a fitting algorithm for $\mathcal{L}(O)$. Then \mathcal{A} is not a sample-efficient PAC learning algorithm, if:*

1. $O \neq \{\exists, \sqcup\}$ and \mathcal{A} always returns a most specific fitting if one exists;
2. $O \neq \{\forall, \sqcap\}$ and \mathcal{A} always returns a most general fitting if one exists;
3. \mathcal{A} always returns a fitting of minimal quantifier depth if some fitting exists.

It is interesting to note that Lemma 1 plays a central role here as well. Indeed, the dualization operation preserves quantifier depth and turns most specific fittings into most general ones, and vice versa. This means that it suffices to show only one of Point 1 or Point 2; the other one follows by dualization. Theorem 3 was known for $\mathcal{L}(\{\sqcap, \exists\}) = \mathcal{EL}$ [9]. We begin with Point 3 and note that it suffices to show Lemma 2 below and then apply Lemma 1.

Lemma 2. *For any O with $\{\exists, \sqcup\} \subseteq O \subseteq O_{all}$, any fitting algorithm for $\mathcal{L}(O)$ that always returns a fitting of minimal quantifier depth if some fitting exists is not a sample-efficient PAC learning algorithm.*

Proof. Assume there is a sample-efficient PAC learning algorithm \mathcal{A} with associated polynomial m that always returns a fitting $\mathcal{L}(O)$ concept of minimal quantifier depth, if some fitting exists. We choose $\varepsilon = 0.4$, $\delta = 0.5$, $\Sigma = \{r, s, t\}$, and a target concept $C_T = \exists t^{n+1}. \top$, for some n to be determined later. We consider the following interpretations $\mathcal{I}_w, \mathcal{J}_w$, for $w \in \{r, s\}^n$:

- for each $w \in \{r, s\}^n$, \mathcal{I}_w is a path of length n whose edges are labeled according to w and whose start element is called a_w ;
- for each $w \in \{r, s\}^n$, \mathcal{J}_w is obtained from \mathcal{I}_w by adding a t -path of length $n + 1$ starting from a_w .

Clearly, each (\mathcal{J}_w, a_w) is a positive example, while each (\mathcal{I}_w, a_w) is a negative example. Let \mathbb{P} be the probability distribution that assigns probability $\frac{1}{2^{n+1}}$ to each of the 2^{n+1} examples $(\mathcal{I}_w, a_w), (\mathcal{J}_w, a_w)$, and 0 to all others. We choose n

large enough so that the probability that a sample of size $m(\frac{1}{\delta}, \frac{1}{\varepsilon}, 3, 2n+1, n+2)$ does not contain both (\mathcal{I}_w, a_w) and (\mathcal{J}_w, a_w) for some $w \in \{r, s\}^n$ is at least $1 - \delta$.

Let P, N be a sample of size $m(\frac{1}{\delta}, \frac{1}{\varepsilon}, 3, 2n+1, n+2)$ and suppose that it satisfies the mentioned property: for each $w \in \{r, s\}^n$, not both $(\mathcal{I}_w, a_w) \in N$ and $(\mathcal{J}_w, a_w) \in P$. Consider the $\mathcal{L}(\{\exists, \sqcup\})$ concept

$$C_0 = \bigsqcup_{(\mathcal{I}_w, a_w) \in P} C_w$$

where C_w is the concept $\exists t_1 \dots \exists t_n. \top$ if $w = t_1 \dots t_n$. Clearly, C_0 has quantifier depth n and fits P, N . Hence, if \mathcal{A} runs on P, N , it returns a fitting $\mathcal{L}(O)$ concept C quantifier depth at most n . However, such C cannot distinguish (\mathcal{I}_w, a_w) and (\mathcal{J}_w, a_w) , for any $w \in \{r, s\}^n$, hence $\text{error}_{\mathbb{P}}(C_T, C) = 0.5 > \varepsilon$. \square

We now turn our attention to Points 1 and 2 of Theorem 3. Inspecting the proof of Point 1 for \mathcal{EL} from [10, Theorem 6] reveals that it actually implies Point 1 for each fragment of \mathcal{ALC} that contains \mathcal{EL} . Hence, it suffices to show Point 2 for $\{\exists, \sqcup\}$, $\{\exists, \sqcup, \sqcap\}$, and $\{\exists, \forall, \sqcap\}$ (note, that adding negation to any of these operator sets leads to full \mathcal{ALC}). We show the first here and the latter two in the appendix. Both are by non-trivial extensions of the proof for the same fact for \mathcal{EL} .

Lemma 3. *Any fitting algorithm for $\mathcal{L}(\{\exists, \sqcup\})$ that always returns a most general fitting concept if one exists is not sample-efficient.*

Proof. Assume towards a contradiction that there is a fitting algorithm \mathcal{A} for $\mathcal{L}(\{\exists, \sqcup\})$ that always returns a most general fitting $\mathcal{L}(\{\exists, \sqcup\})$ concept and which moreover is sample-efficient with associated polynomial m . We choose $\varepsilon = \frac{1}{7}$, $\delta = \frac{8}{9}$, $\Sigma = \{r, s, A\}$, $C_T = A$, and n large enough such that $2^{n-1} > m(\frac{1}{\varepsilon}, \frac{1}{\delta}, 3, 2n, 1)$. We use the following interpretations \mathcal{I}, \mathcal{J} and \mathcal{I}_w for $w \in \{r, s\}^n$:

- \mathcal{I} consists of elements a_1, \dots, a_n satisfying $a_1 \in A^{\mathcal{I}}, a_n \in A^{\mathcal{I}}$, and $(a_i, a_{i+1}) \in r^{\mathcal{I}}$ and $(a_i, a_{i+1}) \in s^{\mathcal{I}}$, for all $i < n$.
- \mathcal{J} consists of elements b_1, \dots, b_n, b_{n+1} satisfying $b_i \in A^{\mathcal{J}}$ for all $i \in \{2, \dots, n-1, n+1\}$ and $(b_i, b_{i+1}) \in r^{\mathcal{I}}, (b_i, b_{i+1}) \in s^{\mathcal{I}}$, for all $i \leq n$, and $(b_{n+1}, b_{n+1}) \in r^{\mathcal{J}}$ and $(b_{n+1}, b_{n+1}) \in s^{\mathcal{J}}$.
- For $w \in \{r, s\}^n$, \mathcal{I}_w is a path of length n whose edges are labeled according to w , whose end satisfies A , and whose start element is denoted a_w .

Let \mathbb{P} be the probability distribution that assigns $\mathbb{P}(\mathcal{I}, a_1) = \mathbb{P}(\mathcal{J}, b_1) = \frac{1}{3}$, $\mathbb{P}(\mathcal{I}_w, a_w) = \frac{1}{3 \cdot 2^n}$ for $w \in \{r, s\}^n$ and probability 0 to all other examples. Note that (\mathcal{I}, a_1) is a positive example, and (\mathcal{J}, b_1) and all (\mathcal{I}_w, a_w) are negative examples for C_T .

Let P, N be a sample of size $m(\frac{1}{\varepsilon}, \frac{1}{\delta}, 3, 2n, 1)$ drawn according to \mathbb{P} . Then with probability at least $1 - \delta$, both $(\mathcal{I}, a_1) \in P$ and $(\mathcal{J}, b_1) \in N$. We argue in the appendix that $C_0 = A \sqcup \bigsqcup_{(\mathcal{I}_w, a_w) \notin N} C_w$ is the unique most general $\mathcal{L}(\{\exists, \sqcup\})$ concept fitting P, N , where C_w is $\exists t_1 \dots \exists t_n. A$ in case $w = t_1 \dots t_n$. Thus, \mathcal{A}

returns (a concept equivalent to) C_0 when run on P, N . However, all negative examples $(\mathcal{I}_w, a_w) \notin N$ are labeled incorrectly by C_0 . The choice of n implies that there are at least 2^{n-1} such examples; thus the error is at least $\frac{1}{6} > \varepsilon$. \square

It remains to discuss the exceptions mentioned in Points 1 and 2 of Theorem 3. Interestingly, bounded fitting is a counter-example for these cases:

Lemma 4. *Bounded fitting for $\mathcal{L}(\{\exists, \sqcup\})$ (resp., $\mathcal{L}(\{\forall, \sqcap\})$) returns a most specific $\mathcal{L}(\{\exists, \sqcup\})$ (resp. most general $\mathcal{L}(\{\forall, \sqcap\})$) concept, if any fitting exists.*

5 SAT-Based Bounded Fitting

Motivated by the NP-completeness of size-restricted fitting established in Theorem 1, we implemented bounded fitting for any logic $\mathcal{L}(O)$ for $O \subseteq O_{\text{all}}$ using a SAT solver to efficiently search for fitting $\mathcal{L}(O)$ concepts of a given size. In this section, we first give an overview of the reduction to SAT and describe two optimizations that improve the performance of the basic reduction. Then, we propose a variant of bounded fitting tailored towards approximate fitting. Finally, we compare our implementation with other concept learning implementations.

We describe the reduction for \mathcal{ALC} . Let P, N be sets of examples and $k \geq 1$. We aim to produce a propositional formula φ that is satisfiable if and only if there is an \mathcal{ALC} concept C that fits P, N with $\|C\| \leq k$. We assume w.l.o.g. that all examples in P and N share a single (finite) interpretation \mathcal{I} . Let Σ_C and Σ_R be the set of concept and role names, respectively, that occur in \mathcal{I} .

The formula φ is the conjunction of two formulas $\varphi_1 \wedge \varphi_2$ whose purpose is to encode syntax trees of \mathcal{ALC} concepts and to evaluate the encoded concepts over \mathcal{I} , respectively. For this task, φ_1 uses variables $x_{i,v}$, $y_{1,i,j}$, and $y_{2,i,j}$ for $i, j \in \{1, \dots, k\}$ and v a node label from $\mathcal{V} = \{\top, \perp, \neg, \sqcap, \sqcup\} \cup \Sigma_C \cup \{\exists r, \forall r \mid r \in \Sigma_R\}$ with the following intended meaning:

- $x_{i,v}$ is true if and only if node i of the syntax tree is labeled with v ,
- $y_{1,i,j}$ is true if and only if node i has the single successor j , and
- $y_{2,i,j}$ is true if and only if node i has the two successors j and $j + 1$.

We treat node 1 as the root of the syntax tree. The clauses in φ_1 enforce that a satisfying assignment to these variables represents a well-formed syntax tree of an \mathcal{ALC} concept. For example, to ensure that nodes that are labeled by a concept name do not have successors and that nodes labeled with \sqcap have two successors, respectively, the clauses

$$x_{i,A} \rightarrow (\neg y_{1,i,j} \wedge \neg y_{2,i,j}) \quad \text{and} \quad x_{i,\sqcap} \rightarrow \bigvee_{i < \ell < k} y_{2,i,\ell}$$

are included in φ_1 for each $i, j \in \{1, \dots, k\}$ with $i < j$ and $A \in \Sigma_C$. Similar clauses encode the rules for nodes labeled with other labels from \mathcal{V} . In case we are interested in a fragment $\mathcal{L}(O)$ of \mathcal{ALC} , we restrict \mathcal{V} accordingly.

To evaluate the encoded \mathcal{ALC} concept, φ_2 uses variables $z_{i,a}$ for $i \in \{1, \dots, k\}$, $a \in \Delta^{\mathcal{I}}$ with the intended meaning that $z_{i,a}$ is true if and only if $a \in C^{\mathcal{I}}$ where C

is the concept represented by node i in the syntax tree. For example, to encode the semantics of concept names and of conjunction \sqcap , φ_2 contains the clauses

$$x_{i,A} \rightarrow z_{i,b}, \quad x_{i,A} \rightarrow \neg z_{i,c}, \quad x_{i,\sqcap} \wedge y_{2,i,j} \rightarrow (z_{i,a} \leftrightarrow (z_{j,a} \wedge z_{j+1,a})) \quad (4)$$

for all $i, j \in \{1, \dots, k\}$, $A \in \Sigma_C$, $b \in A^{\mathcal{I}}$, $c \in \Delta^{\mathcal{I}} \setminus A^{\mathcal{I}}$, and $a \in \Delta^{\mathcal{I}}$.

Finally, the fitting condition is enforced by the following formula:

$$\bigwedge_{(\mathcal{I},a) \in P} z_{1,a} \wedge \bigwedge_{(\mathcal{I},b) \in N} \neg z_{1,b} \quad (5)$$

Given the provided intuitions, it should be clear that we can read off a witness concept from a satisfying assignment of φ . Overall, the constructed formula φ uses $O(k^2 + k|\Sigma| + kn)$ variables, $O(nk^3|\Sigma|)$ clauses, and is of total size $O(nk^3|\Sigma|)$ where $\Sigma = \Sigma_C \cup \Sigma_R$ and $n = \|\mathcal{I}\|$.

5.1 Optimizations

We discuss two optimizations that we implemented on top of the basic reduction described above. Similar optimizations have been implemented in [9, 33].

Taking Advantage of Types The encoding of the semantics of concept names given in (4) produces $k \times |\Delta^{\mathcal{I}}| \times |\Sigma_C|$ clauses. In realistic data sets, both $|\Delta^{\mathcal{I}}|$ and $|\Sigma_C|$ tend to be large. For example, in the Carcinogenesis data set of the SML-Benchmark [39], there are 19 221 individuals and 133 concept names, resulting in more than 10 million clauses for $k = 4$, which is the majority of clauses in the encoding. To reduce the number of clauses, one can exploit that not all combinations of concept names actually occur in realistic data sets. In *carcinogenesis*, only 105 different combinations occur. This allows us to encode the same semantics with fewer clauses, at the cost of introducing additional variables.

The *type* $\text{type}(a)$ of an individual a in an interpretation \mathcal{I} is the set $\{A \in \Sigma_C \mid a \in A^{\mathcal{I}}\}$. Let T be the set of all types that occur in \mathcal{I} , that is, $T = \{\text{type}(a) \mid a \in \Delta^{\mathcal{I}}\}$. For each $i \in \{1, \dots, k\}$ and $t \in T$, we introduce an additional variable $x_{i,t}$ with the intended meaning that $x_{i,t}$ is true if and only if node i is labeled with a concept name that is contained in t . We ensure this with the clauses

$$x_{i,A} \rightarrow x_{i,t} \quad \text{and} \quad x_{i,B} \rightarrow \neg x_{i,t}$$

for each $t \in T$, $i \in \{1, \dots, k\}$, $A \in t$ and $B \in \Sigma_C \setminus t$. We then replace the clauses that enforce the semantics of concept names with

$$x_{i,\text{type}(a)} \rightarrow z_{i,a} \quad \text{and} \quad \neg x_{i,\text{type}(a)} \rightarrow \neg z_{i,a} \vee \neg \ell_i$$

for all $i \in \{1, \dots, k\}$ and $a \in \Delta^{\mathcal{I}}$, where ℓ_i are additional variables that are true if and only if node i is labeled with a concept name. In the carcinogenesis dataset, the modified encoding produces only $4 \cdot 105 \cdot 133 + 4 \cdot 19\,221 \cdot 2 = 209\,628$ clauses.

Syntax Tree Templates A major problem is that different models of φ_1 can represent (syntax trees of) equivalent concepts. For example, already the arguably simple concept $\exists r.A \sqcap \forall s.B$ is encoded in four different models. This phenomenon is well-known also from other domains and is called *symmetry* in the search space [35]. The problem it brings is that the SAT solver often does significant redundant work by trying symmetric solutions, which slows down the solving process, especially for unsatisfiable formulas. A solution is to *break* symmetries by introducing additional clauses to the encoding that restrict how the nodes of a syntax tree may be ordered. Similar symmetries have been observed in other SAT-based implementations of bounded fitting [9, 32].

By the *topology* of the syntax tree of an \mathcal{ALC} concept, we mean the binary tree that results from removing all node labels. It is clear that isomorphic topologies can represent the same set of concepts (up to commutativity of \sqcap/\sqcup). We thus extend φ by clauses that enforce, for each topology isomorphism equivalence class, a canonical ordering of the nodes in the syntax tree. This is achieved by enumerating all binary trees with k nodes, and through a long disjunction, demanding that the encoded syntax tree must have one of these topologies. A similar approach has been taken in [33], but instead of including the large disjunction in the encoding, there is one call of the SAT solver per topology.

Note that the number of binary trees with k nodes grows exponentially with k . To prevent a negative impact on runtime, we limit this process to topologies with at most 10 nodes, after which the topologies are only considered as *prefixes* of the topology of the encoded syntax tree. The used threshold of 10 was determined empirically; higher values resulted in a slowdown on the benchmarks.

It is worth noting that these clauses do not break all symmetries, especially many related to commutativity remain. Essentially, there will always be a tradeoff between the number of clauses needed to break symmetries and the gain in avoiding symmetric solutions. In the extreme case, breaking *all* symmetries will require many clauses and thus likely not result in performance improvements. In our experience, adding some extra clauses avoiding certain syntactical patterns proved beneficial, e.g., \top or \perp directly below \sqcap or \sqcup , \neg directly below \neg , \exists , \forall , and two \sqcap -successors below a \sqcap -node, and the same for \sqcup .

5.2 Approximation Scheme

In practice, there may not always be a concept that perfectly fits the provided examples. In such cases, it is more realistic to view the fitting problem as an optimization problem. We address this by changing the basic bounded fitting algorithm from Algorithm 1 into the approximation scheme described in Algorithm 2. There, an extra variable m stores the best coverage of examples achieved so far. If an exact fitting concept is found (Lines 4/5), the algorithm returns it. Otherwise, m is increased (Line 7). If for the current size-bound k , m cannot be improved (test in Line 3), the algorithm moves on to the next k . Note that this approximation scheme is an *anytime algorithm* in the sense that, if it is interrupted at some point, it may return the best solution found up to this point.

Input: Positive examples P , negative examples N

```

1  $m \leftarrow 1$ ;
2 for  $k := 1, 2, \dots$  do
3   while there is  $\varphi \in \mathcal{L}$  of size  $k$  that fits  $m$  examples from  $P, N$  do
4     if  $m = |P| + |N|$  then                                     // all examples covered?
5       return  $\varphi$ 
6     else                                                         // try to cover more examples
7       increase  $m$ ;

```

Algorithm 2: Approximation version of Bounded Fitting.

This change in the algorithm is reflected in the interface to the SAT solver as follows. We replace clauses (5), that demand exact fitting, with *cardinality constraints*, that demand that at least m of the literals in (5) be true. We exploit the ability to solve *incrementally* provided by some SAT solvers, since the increase of m can be realized by adding clauses. However, whenever k increases, we have to produce a new propositional formula as before.

5.3 Evaluation

Our implementation uses the PySat library [19, 20] to decide satisfiability of φ with the Glucose SAT solver and, if satisfiable, read off a fitting concept from a model of φ . In the following, we refer with ALC-SAT to the base implementation and with ALC-SAT⁺ to the one with the optimized encoding.

We compared our implementation to other concept learning tools CELOE [6], SPARCEL [37], and EvoLearner [16].³ While these tools also support extensions of \mathcal{ALC} such as number restrictions or data properties, we restricted all tools to \mathcal{ALC} for our evaluation. We did not consider fragments in these first experiments since not all tools support them. We also did not compare our implementation against concept learning tools such as NCES2 [24] and DRILL [13], as these require a pretraining step for each data set, which makes direct performance comparison difficult. We evaluate performance in two setups: in the first, we test the ability to find exact fitting concepts on a new benchmark where fitting concepts are guaranteed to exist. In the second, we evaluate the approximation scheme from Section 5.2 on the standard SML benchmark [39].

For evaluating exact fitting we extracted fragments of the YAGO 4.5 knowledge base [36] by taking all facts talking about a given signature. The *family* fragment is based on signature $\{spouse, children, gender, Male, Female\}$, and the *language* fragment uses the role name *hasLanguage* and several concept names that represent languages. We handcrafted 70 \mathcal{ALC} concepts of varying complexity and used them to obtain positive and negative examples by querying the extracted fragments. The task was then to reverse-engineer the \mathcal{ALC} concept from (subsets

³ We use the version of CELOE from DL-Learner 1.5.0 and the version of EvoLearner from ontolearn 0.8.1. The exact learning benchmarks were run on a MacBook Pro with M3 Pro chip and 36GB RAM. The generalization benchmarks were run on a MacBook Pro with M1 Pro Chip and 32 GB RAM.




Fig. 3: Exact learning runtime

of) the examples. For the family fragment we used concepts with high quantifier depth, for example $\exists c.(\exists c.\exists g.F \sqcap \exists c.\exists g.M)$ asking for all individuals having a child that has both a daughter (gender female) and a son (male); we use also deeper concepts, referring to grand-children and so on. For the language fragment we use shallow concepts that contain Boolean combinations of concept names under existential or universal restrictions, for example, $\exists l.(F \sqcup I \sqcup R) \sqcap \forall l.\neg G$.

We highlight the impact of our optimizations on the encoding in Figure 3 (left). The data points stem from experiments on the 25 benchmarks over the language fragment and compare the time ALC-SAT and ALC-SAT+ needed until an exact fitting was found. The data shows that we find exact fitting concepts about twice as fast with optimizations enabled compared. We conjecture that this largely explained by the type encoding which is beneficial in the presence of many concept names.

We also compared the time of ALC-SAT+ to find a fitting concept to that of EvoLearner, CELOE and SPaRCEL over the 45 benchmarks generated from the family fragment. The graph in Figure 3 (right) displays the number of exact fittings found by each tool within a given runtime. It can be seen that ALC-SAT+ finds exact fittings faster than both CELOE or EvoLearner. However, SPaRCEL achieves better results for smaller values for maximum runtime. Both SPaRCEL and ALC-SAT+ find the same number of exact fittings within the timeframe displayed, however, for 5 benchmarks SPaRCEL was unable to find an exact fitting concept whereas ALC-SAT+ was able to find exact fittings for all benchmarks, at the cost of increased execution time.

We evaluated approximate fitting of ALC-SAT+ and the other tools by measuring accuracy and length of the returned concept on the SML-Benchmark using 10-fold validation. Each tool was allowed to run for 5 minutes. The results are given in Table 1 where for each tool and benchmark the first row shows accuracy while the second row shows concept length, both with standard deviation. The accuracy achieved by ALC-SAT+ is comparable to that of EvoLearner, except for Mammographic. The values for accuracy of CELOE and SPaRCEL, in general, are lower. For SPaRCEL, a possible explanation is that the length of concepts is much

Table 1: Generalization results on SML-Benchmarks

	Carcinogenesis	Hepatitis	Lymphography	Mammographic	Mutagenesis	Nctrer
EvoLearner	0.53 ± 0.18	0.58 ± 0.01	0.81 ± 0.12	0.46 ± 0.00	0.78 ± 0.18	0.6 ± 0.04
	7.2 ± 3.52	$3.2, \pm 1.03$	19.6 ± 6.88	1.7 ± 0.48	3.5 ± 1.35	3.3 ± 0.95
CELOE	0.54 ± 0.01	0.41 ± 0.01	0.82 ± 0.11	0.46 ± 0.0	0.56 ± 0.2	0.6 ± 0.04
	3.8 ± 0.42	4.6 ± 1.26	10.8 ± 0.42	1.7 ± 2.21	6.8 ± 0.63	3.9 ± 0.32
SParCEL	0.54 ± 0.1	n/a	0.74 ± 0.13	0.55 ± 0.02	0.73 ± 0.23	0.46 ± 0.08
	860.1 ± 66.54	n/a	164.3 ± 49.48	178.2 ± 21.6	85.1 ± 9.35	65.4 ± 18.84
ALC-SAT ⁺	0.55 ± 0.2	0.58 ± 0.01	0.8 ± 0.09	0.77 ± 0.05	0.81 ± 0.26	0.63 ± 0.07
	6.4 ± 0.7	4 ± 1.7	9.9 ± 0.32	11.3 ± 2.06	9.4 ± 0.84	11.7 ± 0.95

larger than that of concepts returned by EvoLearner, CELOE and ALC-SAT⁺. Small concepts are known to generalize well whereas large concepts are prone to overfitting. In comparison to SPArCEL, EvoLearner, CELOE, and ALC-SAT⁺ are able to find small fitting concepts. While this is ‘by design’ for ALC-SAT⁺, we conjecture that the heuristics in EvoLearner and CELOE prefer shorter concepts as well. Note, that when searching for approximate fitting concepts, ALC-SAT⁺ returns the smallest concept with the highest accuracy found within the time limit and may thus report larger concepts than EvoLearner or CELOE.

6 Conclusion and Future Work

We investigated the bounded fitting paradigm for learning class expressions formulated in the description logic \mathcal{ALC} and its fragments. We analyzed theoretical properties, and we believe that many of our results can be extended to more expressive description logics. Additionally, we provided a preliminary implementation that uses the power of a SAT solver to find good fittings, and compared it with other concept learning tools with encouraging results. From a theoretical perspective, it would be interesting to study bounded fitting in an open world setting where an ontology provides additional background knowledge.

From a practical perspective, we see three relevant avenues for future work. First, OWL 2 DL supports several practically relevant features that are not currently supported by our implementation, such as inverse roles, number restrictions, and data properties. Hence, one natural next step is to extend our implementation to be able to learn (ideally) $SR\mathcal{OIQ}(\mathcal{D})$ concepts. Second, we would like to understand the syntactic shape of fittings that occur in realistic knowledge bases. While currently the shape of the sought concepts is determined by the allowed logical constructors, it would also be possible to further restrict it by bounding the number of, say, disjunctions. Finally, we would like to improve the speed of our implementation either by providing better CNF encodings or by *parallelization* relying on syntax tree templates similar to [33].

Supplemental Material Statement: The source code of our implementation ALC-SAT^+ and instructions on how to reproduce the benchmarks are available at <https://github.com/SAT-based-Concept-Learning/ALCSAT>.

References

1. OWL 2 web ontology language overview. <https://www.w3.org/TR/2012/REC-owl2-overview-20121211/>, accessed: July 24, 2025
2. Anthony, M., Biggs, N.: Computational Learning Theory: An Introduction. Cambridge Tracts in Theoretical Computer Science, Cambridge University Press (1992)
3. Baader, F., Horrocks, I., Lutz, C., Sattler, U.: An Introduction to Description Logics. Cambridge University Press (2017). <https://doi.org/10.1017/9781139025355>
4. Blumer, A., Ehrenfeucht, A., Haussler, D., Warmuth, M.K.: Learnability and the Vapnik-Chervonenkis dimension. *Journal of the ACM (JACM)* **36**(4), 929–965 (1989). <https://doi.org/10.1145/76359.76371>
5. Bordaï, B., Neider, D., Roy, R.: The complexity of learning LTL, CTL and ATL formulas. In: Proceedings of 42nd International Symposium on Theoretical Aspects of Computer Science, STACS. LIPIcs, vol. 327, pp. 19:1–19:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2025). <https://doi.org/10.4230/LIPICS.STACS.2025.19>
6. Bühmann, L., Lehmann, J., Westphal, P., Bin, S.: DL-Learner – structured machine learning on semantic web data. In: Companion Proceedings of the The Web Conference 2018. pp. 467–471 (2018). <https://doi.org/10.1145/3184558.3186235>
7. Camacho, A., McIlraith, S.A.: Learning interpretable models expressed in linear temporal logic. In: Proceedings of the Twenty-Ninth International Conference on Automated Planning and Scheduling, ICAPS. pp. 621–630. AAAI Press (2019). <https://doi.org/10.1609/icaps.v29i1.3529>
8. ten Cate, B., Funk, M., Jung, J.C., Lutz, C.: Fitting algorithms for conjunctive queries. *SIGMOD Rec.* **52**(4), 6–18 (2023). <https://doi.org/10.1145/3641832.3641834>, <https://doi.org/10.1145/3641832.3641834>
9. ten Cate, B., Funk, M., Jung, J.C., Lutz, C.: SAT-based PAC learning of description logic concepts. In: Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI 2023. pp. 3347–3355. ijcai.org (2023). <https://doi.org/10.24963/IJCAI2023/373>
10. ten Cate, B., Funk, M., Jung, J.C., Lutz, C.: SAT-based PAC learning of description logic concepts. *CoRR* **abs/2305.08511** (2023). <https://doi.org/10.48550/arXiv.2305.08511>
11. ten Cate, B., Funk, M., Jung, J.C., Lutz, C.: On the non-efficient PAC learnability of conjunctive queries. *Information Processing Letters* **183**, 106431 (2024). <https://doi.org/10.1016/J.IPL.2023.106431>
12. Cohen, W.W., Hirsh, H.: The learnability of description logics with equality constraints. *Machine Learning* **17**(2-3), 169–199 (1994). <https://doi.org/10.1007/BF00993470>
13. Demir, C., Ngonga Ngomo, A.: Neuro-symbolic class expression learning. In: Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI 2023, 19th-25th August 2023, Macao, SAR, China. pp. 3624–3632. ijcai.org (2023). <https://doi.org/10.24963/ijcai.2023/403>

14. Funk, M., Jung, J.C., Lutz, C.: Actively learning concept and conjunctive queries under \mathcal{EL}^r -ontologies. In: Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21. pp. 1887–1893 (2021). <https://doi.org/10.24963/ijcai.2021/260>
15. Funk, M., Jung, J.C., Lutz, C., Pulcini, H., Wolter, F.: Learning description logic concepts: When can positive and negative examples be separated? In: Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19. pp. 1682–1688 (2019). <https://doi.org/10.24963/ijcai.2019/233>
16. Heindorf, S., Blübaum, L., Düsterhus, N., Werner, T., Golani, V.N., Demir, C., Ngonga Ngomo, A.: EvoLearner: Learning description logics with evolutionary algorithms. In: WWW '22: The ACM Web Conference 2022. pp. 818–828. ACM (2022). <https://doi.org/10.1145/3485447.3511925>
17. Horrocks, I., Kutz, O., Sattler, U.: The even more irresistible SROIQ. In: Doherty, P., Mylopoulos, J., Welty, C.A. (eds.) Proceedings, Tenth International Conference on Principles of Knowledge Representation and Reasoning, Lake District of the United Kingdom, June 2-5, 2006. pp. 57–67. AAAI Press (2006), <http://www.aaai.org/Library/KR/2006/kr06-009.php>
18. Iannone, L., Palmisano, I., Fanizzi, N.: An algorithm based on counterfactuals for concept learning in the semantic web. Applied Intelligence **26**(2), 139–159 (2007). <https://doi.org/10.1007/S10489-006-0011-5>
19. Ignatiev, A., Morgado, A., Marques-Silva, J.: PySAT: A Python toolkit for prototyping with SAT oracles. In: Theory and Applications of Satisfiability Testing – SAT 2018. pp. 428–437 (2018). https://doi.org/10.1007/978-3-319-94144-8_26
20. Ignatiev, A., Tan, Z.L., Karamanos, C.: Towards universally accessible SAT technology. In: 27th International Conference on Theory and Applications of Satisfiability Testing. pp. 4:1–4:11 (2024). <https://doi.org/10.4230/LIPICS.SAT.2024.16>
21. Kearns, M.J., Valiant, L.G.: Cryptographic limitations on learning Boolean formulae and finite automata. Journal of the ACM (JACM) **41**(1), 67–95 (1994). <https://doi.org/10.1145/174644.174647>
22. Kietz, J.: Some lower bounds for the computational complexity of inductive logic programming. In: Machine Learning: ECML-93. pp. 115–123 (1993). https://doi.org/10.1007/3-540-56602-3_131
23. Kouagou, N.J., Heindorf, S., Demir, C., Ngomo, A.N.: Neural class expression synthesis. In: The Semantic Web - 20th International Conference, ESWC 2023. Lecture Notes in Computer Science, vol. 13870, pp. 209–226. Springer (2023). https://doi.org/10.1007/978-3-031-33455-9_13
24. Kouagou, N.J., Heindorf, S., Demir, C., Ngonga Ngomo, A.: Neural class expression synthesis in *ALCHIQ(D)*. In: Machine Learning and Knowledge Discovery in Databases: Research Track - European Conference, ECML PKDD 2023. Lecture Notes in Computer Science, vol. 14172, pp. 196–212. Springer (2023). https://doi.org/10.1007/978-3-031-43421-1_12
25. Lehmann, J.: Learning OWL Class Expressions, Studies on the Semantic Web, vol. 6. IOS Press (2010). <https://doi.org/10.3233/978-1-61499-340-7-i>
26. Lehmann, J., Fanizzi, N., Böhmann, L., d’Amato, C.: Concept learning. In: Perspectives on Ontology Learning, pp. 71–91. AKA / IOS Press (2014), https://jens-lehmann.org/files/2014/pol_concept_learning.pdf
27. Lehmann, J., Hitzler, P.: Concept learning in description logics using refinement operators. Machine Learning **78**, 203–250 (2010). <https://doi.org/10.1007/s10994-009-5146-2>

28. Martins, D.M.L.: Reverse engineering database queries from examples: State-of-the-art, challenges, and research opportunities. *Information Systems* **83**, 89–100 (2019). <https://doi.org/10.1016/J.IS.2019.03.002>
29. Mascle, C., Fijalkow, N., Lagarde, G.: Learning temporal formulas from examples is hard. *CoRR* **abs/2312.16336** (2023). <https://doi.org/10.48550/ARXIV.2312.16336>
30. Neider, D., Gavran, I.: Learning linear temporal properties. In: Proceedings of 18th Conference on Formal Methods in Computer Aided Design FMCAD. pp. 1–10. IEEE (2018). <https://doi.org/10.23919/FMCAD.2018.8603016>
31. Pitt, L., Valiant, L.G.: Computational limitations on learning from examples. *J. ACM* **35**(4), 965–984 (1988). <https://doi.org/10.1145/48014.63140>
32. Pommellet, A., Stan, D., Scatton, S.: SAT-based learning of computation tree logic. In: Proceedings of 12th International Joint Conference on Automated Reasoning IJCAR. Lecture Notes in Computer Science, vol. 14739, pp. 366–385. Springer (2024). https://doi.org/10.1007/978-3-031-63498-7_22
33. Rienner, H.: Exact synthesis of LTL properties from traces. In: Kazmierski, T.J., von Hanxleden, R., Mak, T.S.T. (eds.) 2019 Forum for Specification and Design Languages, FDL 2019, Southampton, United Kingdom, September 2-4, 2019. pp. 1–6. IEEE (2019). <https://doi.org/10.1109/FDL.2019.8876900>
34. Rizzo, G., Fanizzi, N., d’Amato, C.: Class expression induction as concept space exploration: From DL-FOIL to DL-FOCL. *Future Gener. Comput. Syst.* **108**, 256–272 (2020). <https://doi.org/10.1016/J.FUTURE.2020.02.071>
35. Sakallah, K.A.: Symmetry and satisfiability. In: Handbook of Satisfiability - Second Edition, Frontiers in Artificial Intelligence and Applications, vol. 336, pp. 509–570. IOS Press (2021). <https://doi.org/10.3233/FAIA200996>
36. Suchanek, F.M., Alam, M., Bonald, T., Chen, L., Paris, P., Soria, J.: YAGO 4.5: A large and clean knowledge base with a rich taxonomy. In: Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2024. pp. 131–140. ACM (2024). <https://doi.org/10.1145/3626772.3657876>
37. Tran, A.C., Dietrich, J., Guesgen, H.W., Marsland, S.: Parallel symmetric class expression learning. *Journal of Machine Learning Research* **18**, 64:1–64:34 (2017), <https://jmlr.org/papers/v18/14-317.html>
38. Valiant, L.G.: A theory of the learnable. *Communications of the ACM* **27**(11), 1134–1142 (1984). <https://doi.org/10.1145/1968.1972>
39. Westphal, P., Böhmann, L., Bin, S., Jabeen, H., Lehmann, J.: SML-bench - A benchmarking framework for structured machine learning. *Semantic Web* **10**(2), 231–245 (2019). <https://doi.org/10.3233/SW-180308>
40. Zloof, M.M.: Query by example. In: American Federation of Information Processing Societies: 1975 National Computer Conference, 19-22 May 1975, Anaheim, CA, USA. AFIPS Conference Proceedings, vol. 44, pp. 431–438. AFIPS Press (1975). <https://doi.org/10.1145/1499949.1500034>

A Proofs for Section 2

Lemma 5. *For all signatures Σ , \mathcal{ALC} concepts C with signature contained in Σ , interpretations \mathcal{I} , and $a \in \Delta^{\mathcal{I}}$:*

$$a \in C^{\mathcal{I}} \Leftrightarrow a \notin \overline{C}^{\overline{\mathcal{I}}_{\Sigma}}.$$

Proof. We show the lemma by structural induction.

Base Cases. Let \mathcal{I} be an interpretation and $a \in \Delta^{\mathcal{I}}$. Concepts in the base case have the form \top, \perp or A for a concept name $A \in \Sigma \cap \mathbf{N}_{\mathcal{C}}$. In the case of \perp or \top , $a \in \perp^{\mathcal{I}} \Leftrightarrow a \notin \top^{\overline{\mathcal{I}}_{\Sigma}}$ holds because of $\perp^{\mathcal{J}} = \emptyset$ and $\top^{\mathcal{J}} = \Delta^{\mathcal{J}}$ for any interpretation \mathcal{J} . In the case of a concept name A , $a \in A^{\mathcal{I}} \Leftrightarrow a \notin \overline{A}^{\overline{\mathcal{I}}_{\Sigma}}$ follows from the definition of $\overline{\cdot}$.

Induction Step. We distinguish cases:

1. $C = \neg D$:

$$\begin{aligned} & a \in (\neg D)^{\mathcal{I}} \\ \Leftrightarrow & a \notin D^{\mathcal{I}} \\ \stackrel{\text{(I,H)}}{\Leftrightarrow} & a \in \overline{D}^{\overline{\mathcal{I}}_{\Sigma}} \\ \Leftrightarrow & a \notin (\overline{\neg D})^{\overline{\mathcal{I}}_{\Sigma}} = \overline{(\neg D)}^{\overline{\mathcal{I}}_{\Sigma}} \end{aligned}$$

2. $C = (C_1 \sqcap C_2)$:

$$\begin{aligned} & a \in (C_1 \sqcap C_2)^{\mathcal{I}} \\ \Leftrightarrow & a \in C_1^{\mathcal{I}} \cap C_2^{\mathcal{I}} \\ \stackrel{\text{(I,H)}}{\Leftrightarrow} & a \notin \overline{C_1}^{\overline{\mathcal{I}}_{\Sigma}} \text{ and } a \notin \overline{C_2}^{\overline{\mathcal{I}}_{\Sigma}} \\ \Leftrightarrow & a \notin (\overline{C_1}^{\overline{\mathcal{I}}_{\Sigma}} \cup \overline{C_2}^{\overline{\mathcal{I}}_{\Sigma}}) \\ \Leftrightarrow & a \notin (\overline{C_1 \sqcup C_2})^{\overline{\mathcal{I}}_{\Sigma}} = \overline{(C_1 \sqcap C_2)}^{\overline{\mathcal{I}}_{\Sigma}} \end{aligned}$$

3. $C = (C_1 \sqcup C_2)$:

$$\begin{aligned} & a \in (C_1 \sqcup C_2)^{\mathcal{I}} \\ \Leftrightarrow & a \in C_1^{\mathcal{I}} \cup C_2^{\mathcal{I}} \\ \stackrel{\text{(I,H)}}{\Leftrightarrow} & a \notin \overline{C_1}^{\overline{\mathcal{I}}_{\Sigma}} \text{ or } a \notin \overline{C_2}^{\overline{\mathcal{I}}_{\Sigma}} \\ \Leftrightarrow & a \notin (\overline{C_1}^{\overline{\mathcal{I}}_{\Sigma}} \cap \overline{C_2}^{\overline{\mathcal{I}}_{\Sigma}}) \\ \Leftrightarrow & a \notin (\overline{C_1 \sqcap C_2})^{\overline{\mathcal{I}}_{\Sigma}} = \overline{(C_1 \sqcup C_2)}^{\overline{\mathcal{I}}_{\Sigma}} \end{aligned}$$

4. $C = \exists r.C_1$:

$$\begin{aligned} & a \in (\exists r.C_1)^{\mathcal{I}} \\ \Leftrightarrow & \text{there is } (a, b) \in r^{\mathcal{I}} \text{ with } b \in C_1^{\mathcal{I}} \\ \stackrel{\text{(I,H)}}{\Leftrightarrow} & \text{there is } (a, b) \in r^{\overline{\mathcal{I}}_{\Sigma}} \text{ and } b \notin \overline{C_1}^{\overline{\mathcal{I}}_{\Sigma}} \\ \Leftrightarrow & a \notin (\forall r.\overline{C_1})^{\overline{\mathcal{I}}_{\Sigma}} = \overline{(\exists r.C_1)}^{\overline{\mathcal{I}}_{\Sigma}} \end{aligned}$$

5. $C = \forall r.C_1$:

$$\begin{aligned}
 & a \in (\forall r.C_1)^{\mathcal{I}} \\
 \Leftrightarrow & \text{ for all } (a, b) \in r^{\mathcal{I}} : b \in C_1^{\mathcal{I}} \\
 \stackrel{(I,H)}{\Leftrightarrow} & \text{ for all } (a, b) \in r^{\mathcal{I}} : b \notin \overline{C_1}^{\overline{\mathcal{I}}_{\Sigma}} \\
 \Leftrightarrow & \text{ for all } (a, b) \in r^{\mathcal{I}} : b \in (\neg \overline{C_1})^{\overline{\mathcal{I}}_{\Sigma}} \\
 \Leftrightarrow & a \in (\forall r. \neg \overline{C_1})^{\overline{\mathcal{I}}_{\Sigma}} = (\neg \exists r. \overline{C_1})^{\overline{\mathcal{I}}_{\Sigma}} = (\neg \forall r. C_1)^{\overline{\mathcal{I}}_{\Sigma}} \\
 \Leftrightarrow & a \notin \overline{\forall r. C_1}^{\overline{\mathcal{I}}_{\Sigma}}
 \end{aligned}$$

Lemma 1. *For any signature Σ , all sets of examples P and N and \mathcal{ALC} concepts C with signature contained in Σ , C fits P, N if and only if \overline{C} fits $\overline{N}_{\Sigma}, \overline{P}_{\Sigma}$.*

Proof. We show only the "only if"-direction; the "if"-direction is symmetric. Suppose C fits P, N . By definition, $a \in C^{\mathcal{I}}$ for all $(\mathcal{I}, a) \in P$ and $b \notin C^{\mathcal{J}}$ for all $(\mathcal{J}, b) \in N$. By Lemma 5, $a \notin \overline{C}^{\overline{\mathcal{I}}_{\Sigma}}$ and $b \in \overline{C}^{\overline{\mathcal{J}}_{\Sigma}}$ for all $(\mathcal{I}, a) \in P$ and $(\mathcal{J}, b) \in N$. Thus, \overline{C} fits $\overline{N}_{\Sigma}, \overline{P}_{\Sigma}$.

B Proofs for Section 3

Proposition 1. *Size-restricted fitting for $\mathcal{L}(O)$ is NP-hard for any $O \subseteq O_{all}$ with $\exists \in O$. This already holds if only a single positive and a single negative example are allowed, and over a signature consisting of two role names and one concept name.*

We give here the remaining details of the proof for Proposition 1 sketched in the main part. We start with giving a precise definition of interpretations \mathcal{I}, \mathcal{J} .

Recall that $(\{S_1, \dots, S_m\}, k)$ is the hitting set instance we consider, and that we assume that $\bigcup_i S_i = \{1, \dots, n\}$. We first define an interpretation \mathcal{J}' that encodes only the sets S_1, \dots, S_m and then extend it to obtain \mathcal{I} and \mathcal{J} . As described in the main part, \mathcal{J}' contains an r -path of length n for each set S_j , with suitable detours via s . The elements on the r -paths are called $b_{j,0}, \dots, b_{j,n}$ for each set S_j , and the extra elements are called $b'_{j,i}$. Recall that there is an additional sink element c . The interpretation \mathcal{J}' is given as follows.

$$\begin{aligned}
 \Delta^{\mathcal{J}'} &= \{b_{j,i} \mid 1 \leq j \leq m, 0 \leq i \leq n\} \cup \\
 &\quad \{b'_{j,i} \mid 1 \leq j \leq m, 0 \leq i \leq n, i \notin S_j\} \cup \{c\} \\
 A^{\mathcal{J}'} &= \{b_{j,n} \mid 1 \leq j \leq m\} \\
 r^{\mathcal{J}'} &= \{(b_{j,i-1}, b_{j,i}) \mid 1 \leq j \leq m, 1 \leq i \leq n\} \cup \{(c, c)\} \cup \\
 &\quad \{(b'_{j,i}, c) \mid 1 \leq j \leq m, 1 \leq i \leq n, i \notin S_j\} \cup \\
 &\quad \{(b_{j,n}, c) \mid 1 \leq j \leq m\} \\
 s^{\mathcal{J}'} &= \{(b_{j,i-1}, b'_{j,i}), (b'_{j,i}, b_{j,i}) \mid 1 \leq j \leq m, 1 \leq i \leq n, i \notin S_j\} \cup
 \end{aligned}$$

$$\begin{aligned} & \{(b_{j,i-1}, c) \mid 1 \leq j \leq m, 1 \leq i \leq n, i \in S_j\} \cup \{(c, c)\} \cup \\ & \{(b_{j,n}, c) \mid 1 \leq j \leq m\} \end{aligned}$$

We construct \mathcal{I} by extending \mathcal{J}' as follows.

$$\begin{aligned} \Delta^{\mathcal{I}} &= \{a\} \cup \{a_i, a'_i \mid 0 \leq i \leq n\} \cup \Delta^{\mathcal{J}'} \\ A^{\mathcal{I}} &= \{a_n\} \cup A^{\mathcal{J}'} \\ r^{\mathcal{I}} &= \{(a, a_0)\} \cup \{(a_{i-1}, a_i) \mid 1 \leq i \leq n\} \cup r^{\mathcal{J}'} \cup \\ & \quad \{(a, b_{j,0}) \mid 1 \leq j \leq m\} \cup \\ & \quad \{(a'_i, c) \mid 1 \leq i \leq n\} \cup \{(a_n, c)\} \\ s^{\mathcal{I}} &= \{(a_{i-1}, a'_i), (a'_i, a_i) \mid 1 \leq i \leq n\} \cup \{(a_n, c)\} \cup s^{\mathcal{J}'} \end{aligned}$$

The interpretation \mathcal{J} is the extension of \mathcal{J}' with an element b that connects to the starting points of the components representing the sets in S . It is formally defined as follows.

$$\begin{aligned} \Delta^{\mathcal{J}} &= \{b\} \cup \Delta^{\mathcal{J}'} \\ A^{\mathcal{J}} &= A^{\mathcal{J}'} \\ r^{\mathcal{J}} &= \{(b, b_{j,0}) \mid 1 \leq j \leq m\} \cup r^{\mathcal{J}'} \\ s^{\mathcal{J}} &= s^{\mathcal{J}'} \end{aligned}$$

We define $P = \{(\mathcal{I}, a)\}$, $N = \{(\mathcal{J}, b)\}$, and $k' = n + k + 2$. To show Equivalence (2) from the main part, it suffices to show that the following are equivalent.

- (i) S has a hitting set of size at most k ;
- (ii) there is an $\mathcal{L}(\{\exists\})$ concept fitting P, N of size at most k' ;
- (iii) there is an \mathcal{ALC} concept fitting P, N of size at most k' .

For implication (i) \Rightarrow (ii), consider a hitting set H and the concept $D = \exists r.C_n$ defined based on H . Based on the construction of C_n , it is easy to verify that $a \in D^{\mathcal{I}}$. To see that $b \notin D^{\mathcal{J}}$, suppose, with the aim of showing a contradiction, that $b_{j,0} \in C_n^{\mathcal{J}}$, for some j with $1 \leq j \leq m$. Based on the definition of \mathcal{J} and the fact that each C_i needs a reachable point satisfying A , we can verify using induction that:

$$b_{j,\ell} \in C_{n-\ell}^{\mathcal{J}} \text{ for all } \ell \in \{0, \dots, n\}. \quad (6)$$

Consider any element $i \in S_j \cap H$ and consider $b_{j,i-1}$, which by Equation (6) satisfies $b_{j,i-1} \in C_{n-i+1}^{\mathcal{J}}$. By definition of C_{n-i+1} and $i \in H$, C_{n-i+1} is of shape $\exists s.\exists s.C_{n-i}$. By definition of \mathcal{J} , the only s -successor of $b_{j,i-1}$ is c which does not satisfy C_{n-i} , a contradiction.

Implication (ii) \Rightarrow (iii) is trivial.

To prove (iii) \Rightarrow (i), let C_0 be an \mathcal{ALC} concept of size $\|C_0\| \leq k'$ that fits P, N . For now, we assume C_0 to be in negation normal form, that is, negation occurs only in front of concept names. We argue that this is without loss of generality

later. We will show that C_0 can be gradually transformed into an $\mathcal{L}(\{\exists\})$ concept C of size $\|C\| \leq k'$ that fits P, N .

We start restricting the overall shape.

Claim. There is an \mathcal{ALC} concept C of the form $\exists r.D$ of size at most k' that fits P, N .

Proof of the claim. We first observe that C_0 is not the concept name A as $a \notin A^{\mathcal{I}}$. If C_0 is a disjunction or conjunction, we can choose one of the disjuncts or conjuncts, respectively, as a smaller fitting concept C . (This is due to the fact that there are only single positive and negative examples, respectively.) Because both a and b have only r -successors, C must have the form $\exists r.D$ or $\forall r.D$. We can rule out the latter as $a \in (\forall r.D)^{\mathcal{I}}$ would imply $b \in (\forall r.D)^{\mathcal{J}}$. \square

Next, we will eliminate universal restrictions.

Claim. There is an $\mathcal{L}(\{\exists, \sqcap, \sqcup, \neg\})$ concept of shape $\exists r.D$ and size at most k' that fits P, N .

Proof of the claim. Take the concept C fitting P, N resulting from the previous claim. Since C is of the form $\exists r.D$, the definitions of \mathcal{I}, \mathcal{J} implies that:

- $a_0 \in D^{\mathcal{I}}$, and
- $b_{j,0} \notin D^{\mathcal{J}}$ for all $j \in \{1, \dots, m\}$.

Consider the sub-interpretations of \mathcal{I} rooted in a_0 and the sub-interpretations of \mathcal{J} rooted at $b_{j,0}$ for all j . It is readily seen that in all these sub-interpretations, the roles r and s are interpreted as total and functional relations. It is well known that in such interpretations the extensions of $\exists r.E$ and $\forall r.E$ coincide for all \mathcal{ALC} concepts E . Hence, if we replace in D every sub-concept $\forall r.E$ by $\exists r.E$, the resulting concept D' still satisfies $a_0 \in (D')^{\mathcal{I}}$ and $b_{j,0} \notin (D')^{\mathcal{J}}$ for all $j \in \{1, \dots, m\}$. The concept $\exists r.D'$ is then as required by the claim. \square

Next, we will eliminate disjunctions.

Claim. There is a $\mathcal{L}(\{\exists, \sqcap, \neg\})$ concept of shape $\exists r.D$ and of size at most k' that fits P, N .

Proof of the claim. Let $C = \exists r.D$ be a fitting $\mathcal{L}(\{\exists, \sqcap, \sqcup, \neg\})$ concept that exists due to the previous claim. Using distributivity laws, we can transform the concept D into an equivalent concept $C_1 \sqcup \dots \sqcup C_l$ such that each C_i is an $\mathcal{L}(\{\exists, \sqcap, \neg\})$ concept. One can easily verify that $\|C_i\| \leq \|D\|$ for each $i \in \{1, \dots, l\}$ after the transformation. For at least one $i \in \{1, \dots, l\}$ we must have $a' \in C_i^{\mathcal{I}}$ for some r -successor $a' \in \Delta^{\mathcal{I}}$ of a . We choose the concept $\exists r.C_i$ as a fitting $\mathcal{L}(\{\exists, \sqcap, \neg\})$ concept which is at most as large as C . \square

Finally, we proceed by eliminating conjunctions and negations.

Claim. There is a $\mathcal{L}(\{\exists\})$ concept C of shape $C = \exists r.D$ of size at most k' that fits P, N .

Proof of the claim. Let $C = \exists r.D$ be the $\mathcal{L}(\{\exists, \sqcap, \neg\})$ concept that exists due to the previous claim. We need some auxiliary notation. For a finite sequence $w = r_1 \cdots r_m \in \{r, s\}^*$, we abbreviate with $\exists w.C$ the concept $\exists r_1 \dots \exists r_m.C$. Moreover, for some interpretation \mathcal{I}_0 and $a_0 \in \Delta^{\mathcal{I}_0}$, we let $R_w(\mathcal{I}_0, a_0)$ denote the set of all b such that $(a_0, b) \in r_1^{\mathcal{I}_0} \circ \dots \circ r_m^{\mathcal{I}_0}$, where \circ denotes the composition of binary relations.

We will transform C inductively into concepts of the shape $C_1 = \exists w_1.D_1, C_2 = \exists w_2.D_2, \dots$ such that each C_i fits P, N , is of size at most k' , and

1. $R_{w_i}(\mathcal{I}, a) \cap \{a_0, a'_0, \dots, a_n, a'_n\}$ contains a single element and this element satisfies D_i ;
2. $R_{w_i}(\mathcal{J}, b) \neq \emptyset$ and no element from $R_{w_i}(\mathcal{J}, b)$ satisfies D_i .

In the inductive base, we set $w_1 = r$ and $D_1 = D$. Clearly, Properties 1 and 2 are satisfied by this choice.

For the inductive step, take any $C_i = \exists w_i.D_i$. We distinguish cases on the shape of D_i . The easy cases are the following:

- If $D_i = A$, we are done.
- The case $D_i = \top$ is impossible due to Property 2 above.
- The case $D_i = \neg A$ is also impossible due to Property 2 above.
- If $D_i = \exists t.D'_i$ for $t \in \{r, s\}$, we can just extend w_i with t , that is, $w_{i+1} = w_i t$ and $D_{i+1} = D'_i$.

The final case requires a bit more care. Let $D_i = E_1 \sqcap \dots \sqcap E_m$ be a non-empty conjunction. If some conjunct is \top , we can simply remove it. If some conjunct is $\neg A$, we can remove it due to Property 2. If there are two conjuncts $\exists r.F, \exists r.G$, then we can replace them with $\exists r.(F \sqcap G)$, due to functionality, and the same with s .

Suppose now that some conjunct in D_i is A . Then the element witnessing Property 1 has to be a_n . This means that for every other conjunct $\exists t.F$ in D_i , $t \in \{r, s\}$, F is satisfied in c . However, every $\mathcal{L}(\{\exists, \sqcap, \neg\})$ concept in negation normal form satisfied in c is satisfied in every element in $R_{w_i}(\mathcal{J}, b)$. Hence, we can simply set $D_i = A$ and are done.

Hence, we can assume that D_i is actually of shape $\exists r.F_1 \sqcap \exists s.F_2$. We distinguish cases on the shape of the element d from Property 1.

- If $d = a_\ell$, for some ℓ , then we can replace D_i with $\exists s.\exists s.F_1$. This is clearly satisfied in d . Moreover, in this case, elements in $R_{w_i}(\mathcal{J}, b)$ are either c or $b_{j,k}$, and all of them do not satisfy $\exists s.\exists s.F_1$, by the assumption that they do not satisfy $\exists r.F_1 \sqcap \exists s.F_2$. Thus Properties 1 and 2 are preserved.
- If $d = a'_\ell$, for some ℓ , then we can drop $\exists r.F_1$ since F_1 has to be satisfied in c and every $\mathcal{L}(\exists, \sqcap, \neg)$ concept in negation normal form satisfied in c is satisfied in every element in $R_{w_i}(\mathcal{J}, b)$. This preserves Properties 1 and 2.

This finishes the proof of the Claim. \square

Let now be C a fitting $\mathcal{L}(\{\exists\})$ concept that fits P, N and is of size at most k' . We make two observations on the shape of C , using the notation of the proof of the previous claim.

- C is of shape $\exists r.\exists w.A$ for some w : the only other option for the innermost concept would be \top , but $\exists r.\exists w.\top$ is satisfied in b .
- $w \in (r + ss)^+$ since this is the only way to "reach" a_n in the positive example. Moreover, due to the size bound $\|C\| \leq k' = n + k + 2$, the pattern ss occurs at most k times in w .

Due to the second item above we can write $w = w_1 \dots w_n$ with $w_i \in \{r, ss\}$ for all i and such that $w_i = ss$ at most k times. We read off a hitting set H of size at most k from C by taking

$$H = \{i \in \{1, \dots, n\} \mid w_i = ss\}$$

We have to show $H \cap S_i \neq \emptyset$, for every j . Suppose, with the goal of deriving a contradiction, that $H \cap S_j = \emptyset$. Then, by the construction of \mathcal{J} , we have $b_{j,0} \in (\exists w.A)^{\mathcal{J}}$. Hence, $b \in C^{\mathcal{J}}$ and (\mathcal{J}, b) is not a negative example, contradiction.

We have yet to justify our initial assumption that a fitting concept is given in negation normal form. For that, consider any smallest fitting concept C_0 (not necessarily in negation normal form), that is $\|C\| \geq \|C_0\|$ for all fitting concepts C . Let C'_0 be the negation normal form of C_0 . We show that we either can obtain a fitting $\mathcal{L}(\{\exists\})$ concept as outlined above, or that we can construct a smaller fitting concept than C_0 thus reaching a contradiction. We cannot have $\|C'_0\| < \|C_0\|$ as C_0 is a smallest fitting concept. In the case of $\|C'_0\| = \|C_0\|$ we can obtain a fitting $\mathcal{L}(\{\exists\})$ concept from C'_0 as argued above. Now, assume $\|C'_0\| > \|C_0\|$. In this case, C_0 must contain at least one negation and C'_0 must contain at least two negations. A closer look at the above proof yields that from such a concept C'_0 we can obtain a fitting concept C' that is smaller than C_0 , i.e. $\|C'\| < \|C_0\|$. In all transformations on the original fitting concept we either directly choose a subconcept of the original fitting concept or we remove negated concept names to obtain a fitting concept of the desired form. Thus, when an $\mathcal{L}(\{\exists, \sqcap, \neg\})$ concept is reached, negated concept names may have been removed but none have been added. The remaining negated concept names may then be removed as argued above. In conclusion, we construct a fitting concept C by eliminating all negated concept names while possibly eliminating other subconcepts as well. Since C_0 contains at least one negation and C contains none, we have obtained a fitting concept smaller than C_0 contradicting our assumption that C_0 is a smallest fitting concept. It follows that the negation normal form of any smallest fitting concept for these examples is of the same size as the smallest fitting concept, and thus we obtain a fitting concept $\mathcal{L}(\{\exists\})$ by assuming a smallest fitting concept to be given in NNF. \square

Corollary 1 implies that Proposition 1 also holds for all fragments $\mathcal{L}(O)$ with $\{\forall\} \subseteq O \subseteq O_{\text{all}}$. Thus, it also completes the proof of Theorem 1.

C Proofs for Section 4

Lemma 6. *The fitting problem for $\mathcal{L}(\{\forall\})$, $\mathcal{L}(\{\forall, \sqcap\})$, and $\mathcal{L}(\{\forall, \exists, \sqcap\})$ is NP-hard.*

Proof. The proof is inspired by the proof of Theorem 3 in [12].

We show this by reduction from SAT. Let φ be a propositional formula in conjunctive normal form with m clauses C_1, \dots, C_m over n variables x_1, \dots, x_n . For $w = w_1 \cdots w_n \in \{0, 1\}^n$ and some formula ψ , we write $w \models \psi$ if the assignment α defined by setting $\alpha(x_i) = w_i$, for all i , satisfies ψ .

From each clause C_i , we can construct in time polynomial in the size of φ a deterministic finite automaton \mathcal{A}_i satisfying:

- $L(\mathcal{A}_i) \subseteq \{0, 1\}^n$, and
- $w \in L(\mathcal{A}_i)$ if and only if $w \models C_i$, for all $w \in \{0, 1\}^n$.

We transform each $\mathcal{A}_i = (Q_i, \delta_i, q_0^i, F_i)$ into an interpretation \mathcal{I}_i using two role names r_0, r_1 and one concept name F as follows:

$$\begin{aligned} \Delta^{\mathcal{I}_i} &= Q_i \\ F^{\mathcal{I}_i} &= F \\ r_0^{\mathcal{I}_i} &= \{(q, q') \mid \delta(q, 0) = q'\} \\ r_1^{\mathcal{I}_i} &= \{(q, q') \mid \delta(q, 1) = q'\} \end{aligned}$$

Let \mathcal{J} be the interpretation with domain $\Delta^{\mathcal{J}} = \{a\}$ and $F^{\mathcal{J}} = \emptyset$ and $r_0^{\mathcal{J}} = r_1^{\mathcal{J}} = \{(a, a)\}$.

Claim. The following are equivalent for $P = \{(\mathcal{I}_i, q_0^i) \mid 1 \leq i \leq m\}$, $N = \{(\mathcal{J}, a)\}$:

- (i) φ is satisfiable;
- (ii) there is an $\mathcal{L}(\{\forall\})$ concept fitting P, N ;
- (iii) there is an $\mathcal{L}(\{\forall, \exists, \sqcap\})$ concept fitting P, N .

Proof of the claim. To show, (i) \Rightarrow (ii), let α be a satisfying assignment of φ . Based on the definition of the \mathcal{I}_i and \mathcal{J} , it is routine to verify that the concept

$$\forall r_{\alpha(x_1)} \dots \forall r_{\alpha(x_n)}. F$$

fits P, N .

(ii) \Rightarrow (iii) is trivial.

For (iii) \Rightarrow (i), suppose D is an $\mathcal{L}(\{\forall, \exists, \sqcap\})$ concept that fits P, N . Let $D' \in \mathcal{L}(\{\forall, \sqcap\})$ be obtained from D by replacing every existential restriction in D by a universal one (of course over the same role name). Since in all involved examples, r_0 and r_1 are interpreted as total functions, D' also fits P, N . Based on the facts that (a) there is only one negative example and (b) roles are interpreted as functions, we can show that conjunction is not necessary in the sense that we can obtain from D' an $\mathcal{L}(\{\forall\})$ concept D'' that fits P, N . By construction of the \mathcal{I}_i and \mathcal{J} , this concept has to take the shape

$$\forall r_{i_1} \dots \forall r_{i_n}. F$$

It is now routine to verify that $w \models \varphi$ for $w = i_1 \cdots i_n$.

This finishes the proof of the claim and, in fact, of the lemma. \square

Corollary 2. *There is no efficient PAC learning algorithm for $\mathcal{L}(\{\forall\})$, $\mathcal{L}(\{\forall, \square\})$, and $\mathcal{L}(\{\forall, \exists, \square\})$, unless $\text{RP} = \text{NP}$.*

Proof. Let \mathfrak{C} be a class of examples. We say that \mathcal{L} has the *polynomial size fitting property over \mathfrak{C}* if any for every sample P, N from this class \mathfrak{C} : if there is an \mathcal{L} concept fitting P, N , there is one of polynomial size in the size of P, N . We also say that \mathcal{L} is *polynomial time evaluable* if given an example (\mathcal{I}, a) and $C \in \mathcal{L}$, we can determine in polynomial time whether $a \in C^{\mathcal{I}}$. It is well-known that all fragments \mathcal{L} of \mathcal{ALC} are polynomial time evaluable.

A classic theorem by Pitt and Valiant states that if \mathcal{L} has the polynomial size fitting property (possibly over a class \mathfrak{C}) and is polynomial time evaluable (over \mathfrak{C}), then efficient PAC learnability implies that the fitting problem (over \mathfrak{C}) is in RP [31].

Let us denote with \mathfrak{C}_0 the class of all interpretations obtained from the construction in the hardness proof in Lemma 6. It is not difficult to see that $\mathcal{L}(\{\forall\})$, $\mathcal{L}(\{\forall, \square\})$, and $\mathcal{L}(\{\forall, \exists, \square\})$ have the polynomial fitting property over the class \mathfrak{C}_0 .

To show the statement from the corollary, let us assume that there is an efficient PAC learning algorithm for \mathcal{L} from $\mathcal{L}(\{\forall\})$, $\mathcal{L}(\{\forall, \square\})$, and $\mathcal{L}(\{\forall, \exists, \square\})$. By Pitt's and Valiant's theorem, the fitting problem for \mathcal{L} over \mathfrak{C}_0 is in RP . The proof of Lemma 6 implies that the fitting problem for \mathcal{L} is NP -hard already over that class \mathfrak{C}_0 . Altogether, this implies $\text{RP} = \text{NP}$. \square

Lemma 2. *For any O with $\{\exists, \square\} \subseteq O \subseteq O_{\text{all}}$, any fitting algorithm for $\mathcal{L}(O)$ that always returns a fitting of minimal quantifier depth if some fitting exists is not a sample-efficient PAC learning algorithm.*

Proof. It remains to argue that we can choose n large enough such that the probability that a sample of size $m(1/\delta, 1/\varepsilon, 3, 2n + 1, n + 2)$ does not contain both (\mathcal{I}_w, a_w) and (\mathcal{J}_w, a_w) for some $w \in \{r, s\}^n$ is at least $1 - \delta$. Since we know that m is a polynomial, there is a polynomial $p(n)$ such that $p(n) = m(1/\delta, 1/\varepsilon, 3, 2n + 1, n + 2)$ for all n .

Let \mathbb{P}'_n be the uniform distribution over a set of 2^n elements. We observe that the probability of seeing in a sample of \mathbb{P} of size $p(n)$ not both (\mathcal{I}_w, a_w) and (\mathcal{J}_w, a_w) for some $w \in \{r, s\}^n$ is greater than the probability of seeing in a sample of size $p(n)$ over \mathbb{P}'_n no element twice.

The probability of sampling ℓ different elements from a uniform distribution over an N -element set is computed as the following fraction:

$$\frac{\prod_{i=0}^{\ell-1} (N - i)}{N^\ell} = \frac{N!}{N^\ell \cdot (N - \ell)!} \tag{7}$$

Hence, it suffices to show that the limit of the fraction in the right of Equation (7) when we put $N = 2^n$ and $\ell = p(n)$ and let n go towards ∞ is 1. This is precisely the statement of [10, Lemma 6]. \square

Lemma 3. *Any fitting algorithm for $\mathcal{L}(\{\exists, \square\})$ that always returns a most general fitting concept if one exists is not sample-efficient.*

Proof. We give the missing argument that $C_0 = A \sqcup \sqcup_{(\mathcal{I}_w, a_w) \notin N} C_w$ is the unique most general $\mathcal{L}(\{\exists, \sqcup\})$ concept fitting P, N . Clearly, C_0 fits P, N . So let D be any $\mathcal{L}(\{\exists, \sqcup\})$ that fits P, N . We have to show that $D \sqsubseteq C_0$.

Consider the syntax trees of some $\mathcal{L}(\{\exists, \sqcup\})$ concept. A *path in the syntax tree* is any sequence $r_1 \dots r_k A$ that can be read by starting at the root of the syntax tree, reading only the role name for $\exists r$ -nodes, skipping \sqcup -nodes, and finally reading the label of the leaf (which is either a concept name or \top).

It can be verified that $D \sqsubseteq C_0$ holds if and only if all paths in the syntax tree of D are paths in the syntax tree of C_0 . So consider a path $r_1 \dots r_k A$ in the syntax tree of D . We call k the length of the path. Since $(\mathcal{J}, b_1) \in N$ and D fits P, N , the path cannot have length $1, \dots, n-1, n+1$ (otherwise, it would incorrectly label (\mathcal{J}, b_1)). Hence, it can only have length 0 or length n . If it is of length 0, then it has to have the shape A and this path is also contained in C_0 , by definition. Otherwise, it is of shape wA with $w \in \{r, s\}^n$. But since D fits P, N , we know that $(\mathcal{I}_w, a_w) \notin N$. By definition, C_0 contains the path wA in this case. \square

It remains to cover the cases $\mathcal{L}(\{\exists, \forall, \sqcap\})$ and $\mathcal{L}(\{\exists, \sqcup, \sqcap\})$ as announced in the body of the paper. Let us state the lemmas formally.

Lemma 7. *Let \mathcal{A} be a fitting algorithm for $\mathcal{L}(\{\exists, \sqcup, \sqcap\})$ that always returns always a most general fitting if one exists. Then \mathcal{A} is not a sample-efficient PAC learning algorithm.*

Lemma 8. *Let \mathcal{A} be a fitting algorithm for $\mathcal{L}(\{\exists, \forall, \sqcap\})$ that always returns always a most general fitting if one exists. Then \mathcal{A} is not a sample-efficient PAC learning algorithm.*

The proofs of both Lemmas 7 and 8 are (non-trivial) variations of the proof for the same fact for $\mathcal{EL} = \mathcal{L}(\{\exists, \sqcap\})$ [10, Theorem 3/Theorem 7]. We require some additional notation and facts which we introduce next.

Every \mathcal{EL} concept C may be viewed as tree-shaped interpretation \mathcal{I}_C in an obvious way, e.g. $C = \exists r. \exists s. A$ results in a path of length 2 which ends in a node satisfying A and whose edges are labeled with r and s in this order. We denote with a_C the root node of this tree. Let $\mathcal{I}_1, \mathcal{I}_2$ be interpretations and Σ a signature. A Σ -simulation from \mathcal{I}_1 to \mathcal{I}_2 is a relation $S \subseteq \Delta^{\mathcal{I}_1} \times \Delta^{\mathcal{I}_2}$ such that for all $(d_1, d_2) \in S$:

1. if $d_1 \in A^{\mathcal{I}_1}$ with $A \in \Sigma$, then $d_2 \in A^{\mathcal{I}_2}$;
2. if $(d_1, e_1) \in r^{\mathcal{I}_1}$ with $r \in \Sigma$, there is $(d_2, e_2) \in r^{\mathcal{I}_2}$ such that $(e_1, e_2) \in S$.

For $d_1 \in \Delta^{\mathcal{I}_1}$ and $d_2 \in \Delta^{\mathcal{I}_2}$, we write $(\mathcal{I}_1, d_1) \preceq_{\Sigma} (\mathcal{I}_2, d_2)$ if there is a Σ -simulation S from \mathcal{I}_1 to \mathcal{I}_2 with $(d_1, d_2) \in S$.

Simulations are important since on the one hand they conveniently characterize evaluation and subsumption in \mathcal{EL} , as defined in the following two well-known statements.

Lemma 9. *For all \mathcal{EL} -concepts C , interpretations \mathcal{I} , and $d \in C^{\mathcal{I}}$, we have $d \in C^{\mathcal{I}}$ iff $(\mathcal{I}_C, a_C) \preceq (\mathcal{I}, d)$.*

Lemma 10. *For all \mathcal{EL} -concepts C, D , we have $C \sqsubseteq D$ iff $(\mathcal{I}_D, a_D) \preceq (\mathcal{I}_C, a_C)$.*

We next introduce the notion of *simulation duals*.

Definition 2. *Let (\mathcal{I}, a) be an example and Σ a signature. A set M of examples is a Σ -simulation dual of (\mathcal{I}, a) if for all examples (\mathcal{I}', a') , the following holds:*

$$(\mathcal{I}, a) \preceq_{\Sigma} (\mathcal{I}', a') \quad \text{iff} \quad (\mathcal{I}', a') \not\preceq_{\Sigma} (\mathcal{I}'', a'') \text{ for all } (\mathcal{I}'', a'') \in M.$$

We will rely on the following theorem from [9].

Theorem 4 ([9, Theorem 3]). *Given an \mathcal{EL} concept C and a signature Σ , a Σ -simulation dual M of (\mathcal{I}_C, a_C) of size $\|M\| \leq 3 \cdot |\Sigma| \cdot \|C\|^2$ can be computed in polynomial time. Moreover, if \mathcal{I}_C , when viewed as a set, contains only a single Σ -assertion that mentions a_C , then M is a singleton.*

We are now ready to move on to the proofs of Lemmas 7 and 8. Both proofs use signature $\Sigma = \{A, B, r\}$ and a family of concept names $(C_n)_{n \geq 0}$ inductively defined as follows:

$$\begin{aligned} C_0 &= \top \\ C_{i+1} &= \exists r.(A \sqcap B \sqcap C_i), \text{ for all } i \geq 0. \end{aligned}$$

We give now first the (easier) proof of Lemma 7, which is rather close to the proof of [10, Theorem 3/7].

Proof (of Lemma 7). Let \mathcal{A} be a fitting algorithm for $\mathcal{L}(\{\exists, \sqcup, \sqcap\})$ that always returns a most general fitting if one exists, and additionally is a sample-efficient PAC learning algorithm with associated polynomial m . Set $\delta = \varepsilon = \frac{1}{2}$ and n large enough so that

$$2^n > 4m(1/\delta, 1/\varepsilon, 3, p(n), 3n)$$

for some polynomial p to be determined below.

As signature we use $\Sigma = \{A, B, r\}$. As target concept C_T we use $C_T := C_n$. We use only negative examples. To construct them, let us first define a set of \mathcal{EL} concepts S_n , inductively as follows:

$$\begin{aligned} S_0 &= \{\top\} \\ S_{i+1} &= \{\exists r.(\alpha \sqcap C) \mid C \in S_i, \alpha \in \{A, B\}\}, \text{ for } i \geq 0 \end{aligned}$$

Clearly, S_n contains 2^n \mathcal{EL} concepts, and $C_T \sqsubseteq C$ and $\|C\| \leq 3n$, for every $C \in S_n$. Consider the example (\mathcal{I}_C, a_C) with $C \in S_n$. Note that \mathcal{I}_C contains a single assertion that mentions a_C . By Theorem 4, (\mathcal{I}_C, a_C) has a singleton Σ -simulation dual $\{(\mathcal{I}'_C, a'_C)\}$ with $\|\mathcal{I}'_C\| \leq 3 \cdot |\Sigma| \cdot 3n$. These duals are negative examples for C_T since $C_T \sqsubseteq C$, using Lemma 10 and the definition of duals.

Let the probability distribution \mathbb{P} assign probability $1/2^n$ to each (\mathcal{I}'_C, a'_C) , $C \in S_n$.

Claim 1. For each $C \in S_n$, C is the most general $\mathcal{L}(\{\exists, \sqcup, \sqcap\})$ concept that fits the negative example (\mathcal{I}'_C, a'_C) .

Proof Claim 1. The very same statement with \mathcal{EL} in place of $\mathcal{L}(\{\exists, \sqcup, \sqcap\})$ is stated in the original proof, and we rely on this fact. Let D be an $\mathcal{L}(\{\exists, \sqcup, \sqcap\})$ concept that fits the negative example (\mathcal{I}'_C, a'_C) . We show that $D \sqsubseteq C$, hence C is also the most general $\mathcal{L}(\{\exists, \sqcup, \sqcap\})$ concept that fits the negative example (\mathcal{I}'_C, a'_C) . We can equivalently write D as a disjunction of \mathcal{EL} -concepts $D \equiv D_1 \sqcup \dots \sqcup D_n$. Since D fits negative example (\mathcal{I}'_C, a'_C) , Lemma 9 yields $(\mathcal{I}_{D_i}, a_{D_i}) \not\leq (\mathcal{I}'_C, a'_C)$, for all i . By definition of duals, we get $(\mathcal{I}_C, a_C) \preceq (\mathcal{I}_{D_i}, a_{D_i})$, for all i . Lemma 10 yields $D_i \sqsubseteq C$, for all i , and hence $D \sqsubseteq C$. \dashv

Let now \emptyset, N be a sample of $m(1/\delta, 1/\varepsilon, 3, p(n), 3n)$ examples from \mathbb{P} (recall that there are no positive examples).

Claim 2. $C_H = \prod_{(\mathcal{I}'_C, a'_C) \in N} C$ is the most general $\mathcal{L}(\{\exists, \forall, \sqcap\})$ concept fitting \emptyset, N .

Proof of Claim 2. We first observe that $C_H = \prod_{(\mathcal{I}'_C, a'_C) \in N} C$ fits \emptyset, N , since, by the definition of duals, (\mathcal{I}'_C, a'_C) is a negative example for each such C .

We now verify that it is the most general one. Let D be any $\mathcal{L}(\{\exists, \forall, \sqcap\})$ concept that fits \emptyset, N . By Claim 1, $D \sqsubseteq C$, for each $(\mathcal{I}'_C, a'_C) \in N$. Hence, $D \sqsubseteq C_H$. \dashv

Hence, (a concept equivalent to) C_H is output by the algorithm. We claim that C_H labels all negative examples not in the sample wrong. Indeed, $C_H \not\sqsubseteq C'$ for any C' with $(\mathcal{I}''_{C'}, a'_{C'}) \notin N$ (not difficult to see and shown in the original proof). By definition of duals, $(\mathcal{I}''_{C'}, a'_{C'})$ is labeled positively by C_H . Since there are 2^n negative examples, each with probability $1/2^n$, the choice of n implies that the error is greater than $\varepsilon = 1/2$. \square

We move to the proof of Lemma 8. It is a variant of the above proof, but needs much more involved constructions.

Proof (of Lemma 8). Let \mathcal{A} be a fitting algorithm for $\mathcal{L}(\{\exists, \forall, \sqcap\})$ that always returns a most general fitting if one exists, and additionally is a sample-efficient PAC learning algorithm with associated polynomial m . Set $\delta = \varepsilon = \frac{1}{2}$ and n large enough so that

$$2^n > 4m\left(\frac{1}{\delta}, \frac{1}{\varepsilon}, 3, p(n), 3n\right)$$

for some polynomial p to be determined below.

As signature we use $\Sigma = \{A, B, r\}$. As target concept C_T we use $C_T := C_n$ which is defined inductively as follows:

$$\begin{aligned} C_0 &= \top \\ C_{i+1} &= \exists r.(A \sqcap B \sqcap C_i), \text{ for all } i \geq 0. \end{aligned}$$

We use a single positive example (\mathcal{I}, a_0) and 2^n negative examples. The mentioned proof of [10, Theorem 3/7] uses the same negative examples, but no positive example. Intuitively, we use the additional positive example to “switch off” universal quantification \forall .

The positive example (\mathcal{I}, a_0) is defined as follows:

$$\begin{aligned}\Delta^{\mathcal{I}} &= \{a_i, b_i, c_i \mid 0 \leq i \leq n\} \\ A^{\mathcal{I}} = B^{\mathcal{I}} &= \{a_1, \dots, a_n\} \\ r^{\mathcal{I}} &= \{(a_i, a_{i+1}), (a_i, b_i), (a_i, c_i), (c_i, c_i) \mid 0 \leq i < n\}\end{aligned}$$

Intuitively, \mathcal{I} is an r -path of length n in which each element (except the last a_n) has one outgoing r -edge to a "dead end" and one outgoing r -edge to an r -self-loop. It is readily verified that there is no $\mathcal{L}(\{\exists, \forall, \sqcup\})$ concept D such that $a_i \in (\forall r.D)^{\mathcal{I}}$, for $1 \leq i < n$. We will rely on this fact below.

To construct the negative examples, let us first define a set of \mathcal{EL} concepts S_n , inductively as follows:

$$\begin{aligned}S_0 &= \{\top\} \\ S_{i+1} &= \{\exists r.(\alpha \sqcap C) \mid C \in S_i, \alpha \in \{A, B\}\}, \text{ for } i \geq 0\end{aligned}$$

Clearly, S_n contains 2^n \mathcal{EL} concepts, and $C_T \sqsubseteq C$ and $\|C\| \leq 3n$, for every $C \in S_n$. Consider the example (\mathcal{I}_C, a_C) with $C \in S_n$. Note that \mathcal{I}_C contains a single assertion that mentions a_C . By Theorem 4, (\mathcal{I}_C, a_C) has a singleton Σ -simulation dual $\{(\mathcal{I}''_C, a'_C)\}$ with $\|\mathcal{I}''_C\| \leq 3 \cdot |\Sigma| \cdot 3n$. These duals are negative examples for C_T since $C_T \sqsubseteq C$, using Lemma 10 and the definition of duals. Clearly, (\mathcal{I}, a_0) is a positive example for C_T . We slightly modify \mathcal{I}''_C by adding a fresh r -successor to every element, calling the resulting interpretation \mathcal{I}'_C . Note that (\mathcal{I}'_C, a'_C) and (\mathcal{I}''_C, a'_C) are simulation-equivalent, due to totality of $r^{\mathcal{I}''_C}$. So they satisfy precisely the same \mathcal{EL} concepts, and hence all (\mathcal{I}'_C, a'_C) are negative examples.

Let the probability distribution \mathbb{P} assign probability $1/2$ to (positive example) (\mathcal{I}, a_0) and probability $1/2^{n+1}$ to each (\mathcal{I}'_C, a'_C) , $C \in S_n$.

Claim 1. For each $C \in S_n$, C is the most general \mathcal{EL} concept that fits the negative example (\mathcal{I}'_C, a'_C) .

Claim 1 is from the original proof, so we do not prove it here. Instead, we prove another claim that allows us to reuse Claim 1. Its proof relies on the positive example.

Claim 2. Let $C \in S_n$ and consider $P = \{(\mathcal{I}, a_0)\}$ and $N = \{(\mathcal{I}'_C, a'_C)\}$. If $D \in \mathcal{L}(\{\exists, \forall, \sqcap\})$ fits P, N , then $D \sqsubseteq C$.

Proof of Claim 2. We inductively show that, for $\ell = 0 \dots n$:

- (*) if there is some $\mathcal{L}(\{\exists, \forall, \sqcap\})$ concept D that distinguishes (\mathcal{I}, a_ℓ) from some (\mathcal{I}'_C, d) , $d \in \Delta^{\mathcal{I}''_C}$, then there is some \mathcal{EL} concept \widehat{D} with $D \sqsubseteq \widehat{D}$ that distinguishes (\mathcal{I}, a_ℓ) from (\mathcal{I}'_C, d) .

For the inductive start, consider a_n . By construction, a_n does not satisfy any $\mathcal{L}(\{\exists, \forall, \sqcap\})$ concept of shape $\forall r.E$. Let now have D a top-level conjunct of shape $\exists r.E$. Clearly, E cannot have top-level conjunct A/B . Moreover, E cannot be of

shape $\forall r.C$ since any such E is satisfied in d , due to the fresh r -successors (added when transitioning from \mathcal{I}''_C to \mathcal{I}'_C). Hence, E has top-level conjunct $\exists r.F$, then only c_n can be the witness of E . But the only $\mathcal{L}(\{\exists, \forall, \sqcap\})$ concepts satisfied by c_n are equivalent to $\exists r^u.\top$ for some $u \geq 0$. But these are satisfied by d as well, so do not help in distinguishing. It follows that D has a top-level conjunct A/B and this is the distinguishing \mathcal{EL} concept that we claimed to exist.

For the inductive step, consider some a_ℓ , and let $D \in \mathcal{L}(\{\exists, \forall, \sqcap\})$ be the concept distinguishing a_ℓ from d of the form

$$D = A_1 \sqcap \dots \sqcap A_\ell \sqcap \forall r.D' \sqcap \exists r.D_1 \sqcap \dots \sqcap \exists r.D_k.$$

We analyze the shape of D . Clearly, by construction of \mathcal{I} and since (\mathcal{I}, a_ℓ) is a positive example for D , we crucially have $D' = \top$ as a_ℓ does not satisfy any non-trivial concept of shape $\forall r.E$. Hence, there is some D_i that is satisfied in a successor of a_ℓ , but in none of the successors of d . We distinguish cases:

- D_i cannot be of shape $D_i = \forall r.D''$, since any such concept is satisfied at the fresh r -successor of d (added when transitioning from \mathcal{I}''_C to \mathcal{I}'_C).
- If D_i contains some top-level conjunct A/B , then the mentioned successor of a_ℓ satisfying D_i , is $a_{\ell+1}$ and we can take $\widehat{D} = \exists r.\widehat{D}_i$ where \widehat{D}_i is the \mathcal{EL} concept obtained from the induction hypothesis applied to $a_{\ell+1}$ and any r -successor of d .
- If D_i contains no top-level conjunct of shape A/B , but a top-level $\exists r.C$, then the mentioned successor of a_ℓ cannot be b_ℓ , since b_ℓ does not have any r -successor. It can also not be c_ℓ since the only $\mathcal{L}(\{\exists, \forall, \sqcap\})$ concepts that are satisfied in c_ℓ are equivalent to $\exists r^u.\top$ for some $u \geq 0$ and each such concept is satisfied in some successor of d , due to totality of $r^{\mathcal{I}''_C}$. Hence, the mentioned successor of a_ℓ satisfying D_i , is $a_{\ell+1}$, and it remains to apply induction as in the previous case.

This finishes the proof of Statement (*). To see the statement from the claim, let $D \in \mathcal{L}(\{\exists, \forall, \sqcap\})$ fit P, N . By (*), there is an \mathcal{EL} concept \widehat{D} that fits P, N and such that $D \sqsubseteq \widehat{D}$. By Claim 1, $\widehat{D} \sqsubseteq C$, hence also $D \sqsubseteq C$ as required. \dashv

Let now P, N be a sample of $m(1/\delta, 1/\varepsilon, 3, p(n), 3n)$ examples sampled from \mathbb{P} . With probability at least $1 - \delta = 1/2$, P contains the (only) positive example (\mathcal{I}, a_0) .

Claim 3. $C_H = \bigcap_{(\mathcal{I}''_C, a'_C) \in N} C$ is the most general $\mathcal{L}(\{\exists, \forall, \sqcap\})$ concept fitting P, N .

Proof of Claim 3. We first verify that $C_H = \bigcap_{(\mathcal{I}''_C, a'_C) \in N} C$ fits P, N . By construction of \mathcal{I} , (\mathcal{I}, a_0) is a positive example for each $C \in S_n$. Moreover, by the definition of duals, (\mathcal{I}''_C, a'_C) is a negative example for each such C .

We now verify that it is the most general one. Let D be any $\mathcal{L}(\{\exists, \forall, \sqcap\})$ concept that fits P, N . By Claim 1 and Claim 2 together, $D \sqsubseteq C$, for each $(\mathcal{I}''_C, a'_C) \in N$. Hence, $D \sqsubseteq C_H$. \dashv

Hence, (a concept equivalent to) C_H is output by the algorithm. We claim that C_H labels all negative examples not in the sample wrong. Indeed, $C_H \not\sqsubseteq C'$ for any C' with $(\mathcal{I}''_{C'}, a'_C) \notin N$ (not difficult to see and shown in the original proof). By definition of duals, $(\mathcal{I}''_{C'}, a'_C)$ is labeled positively by C_H . Since there are 2^n negative examples, each with probability $1/2^{n+1}$, the choice of n implies that the error is greater than $\varepsilon = 1/2$. \square

Lemma 4. *Bounded fitting for $\mathcal{L}(\{\exists, \sqcup\})$ (resp., $\mathcal{L}(\{\forall, \sqcap\})$) returns a most specific $\mathcal{L}(\{\exists, \sqcup\})$ (resp. most general $\mathcal{L}(\{\forall, \sqcap\})$) concept, if any fitting exists.*

Proof. We prove the statement for $\mathcal{L}(\{\exists, \sqcup\})$, the other part is dual, due to Lemma 1. Let P, N be sets of positive and negative examples and let C be an $\mathcal{L}(\{\exists, \sqcup\})$ concept of minimal size fitting P, N . Such concept is returned by bounded fitting. Due to minimality of C , we can assume that

(*) C does not have a subconcept of shape $\exists r.C_1 \sqcup \exists r.C_2$,

as it could be equivalently replaced by the smaller concept $\exists r.(C_1 \sqcup C_2)$.

Let D be any concept at least as specific than C , that is, $D \sqsubseteq C$. Without loss of generality, we can assume that also D satisfies (*). Consider the syntax trees of C and D . A *path in a syntax tree* is any sequence $r_1 \dots r_k A$ that can be read by starting at the root of the syntax tree, reading only the role name for $\exists r$ -nodes, skipping \sqcup -nodes, and finally reading the label A/\top of the leaf.

It can be verified that $D \sqsubseteq C$ holds if and only if all paths in the syntax tree of D are paths in the syntax tree of C . As $D \sqsubseteq C$ by assumption, all paths in D are paths in C . Based on the fact that C and D satisfy Property (*) above, one can then show that there is an injective homomorphism (which is just a mapping that preserves node labels) from the syntax tree of D into the syntax tree of C . But this means that the syntax tree of D is at most as large as the syntax tree of C . Equality is only the case if $C \equiv D$.

Hence, if D would also fit P, N , then bounded fitting would have returned D instead of C , a contradiction. \square