

Foundations for Uniform Interpolation and Forgetting in Expressive Description Logics

Carsten Lutz

Fachbereich Informatik
Universität Bremen, Germany

Frank Wolter

Department of Computer Science
University of Liverpool, UK

Abstract

We study uniform interpolation and forgetting in the description logic \mathcal{ALC} . Our main results are model-theoretic characterizations of uniform interpolants and their existence in terms of bisimulations, tight complexity bounds for deciding the existence of uniform interpolants, an approach to computing interpolants when they exist, and tight bounds on their size. We use a mix of model-theoretic and automata-theoretic methods that, as a by-product, also provides characterizations of and decision procedures for conservative extensions.

1 Introduction

In Description Logic (DL), a *TBox* or *ontology* is a logical theory that describes the conceptual knowledge of an application domain using a set of appropriate predicate symbols. For example, in the domain of universities and students, the predicate symbols could include the concept names *Uni*, *Undergrad*, and *Grad*, and the role name *has.student*. When working with an ontology, it is often useful to eliminate some of the used predicates while retaining the meaning of all remaining ones. For example, when *re-using* an existing ontology in a new application, then typically only a very small fraction of the predicates is of interest. Instead of re-using the whole ontology, one can thus use the potentially much smaller ontology that results from an elimination of the non-relevant predicates. Another reason for eliminating predicates is *predicate hiding*, i.e., an ontology is to be published, but some part of it should be concealed from the public because it is confidential [8]. Finally, one can view the result of predicate elimination as an approach to *ontology summary*: the resulting, smaller and more focussed ontology summarizes what the original ontology says about the remaining predicates.

The idea of eliminating predicates has been studied in AI under the name of *forgetting a signature (set of predicates)* Σ , i.e., rewriting a knowledge base K such that it does not use predicates from Σ anymore and still has the same logical consequences that do not refer to predicates from Σ [13]. In propositional logic, forgetting is also known as *variable elimination* [12]. In mathematical logic, forgetting has been investigated under the dual notion of *uniform interpolation w.r.t. a*

signature Σ , i.e., rewriting a formula φ such that it uses *only* predicates from Σ and has the same logical consequences formulated *only* in Σ . The result of this rewriting is then the *uniform interpolant* of φ w.r.t. Σ . This notion can be seen as a generalization of the more widely known Craig interpolation.

Due to the various applications briefly discussed above, forgetting and uniform interpolation receive increased interest also in a DL context [4; 19; 18; 16; 11; 10]. Here, the knowledge base K resp. formula φ is replaced with a TBox \mathcal{T} . In fact, uniform interpolation is rather well-understood in lightweight DLs such as DL-Lite and \mathcal{EL} : there, uniform interpolants of a TBox \mathcal{T} can often be expressed in the DL in which \mathcal{T} is formulated [11; 10] and practical experiments have confirmed the usefulness and feasibility of their computation [10]. The situation is different for ‘expressive’ DLs such as \mathcal{ALC} and its various extensions, where much less is known. There is a thorough understanding of uniform interpolation on the level of *concepts*, i.e., computing uniform interpolants of concepts instead of TBoxes [14; 18], which is also what the literature on uniform interpolants in modal logic is about [15; 9]. On the TBox level, a basic observation is that there are very simple \mathcal{ALC} -TBoxes and signatures Σ such that the uniform interpolant of \mathcal{T} w.r.t. Σ cannot be expressed in \mathcal{ALC} (nor in first-order predicate logic) [6]. A scheme for approximating (existing or non-existing) interpolants of \mathcal{ALC} -TBoxes was devised in [16]. In [19], an attempt is made to improve this to an algorithm that computes uniform interpolants of \mathcal{ALC} -TBoxes in an exact way, and also decides their existence (resp. expressibility in \mathcal{ALC}). Unfortunately, that algorithm turns out to be incorrect.

The aim of this paper is to lay foundations for uniform interpolation in \mathcal{ALC} and other expressive DLs, with a focus on (i) model-theoretic characterizations of uniform interpolants and their existence; (ii) deciding the existence of uniform interpolants and computing them in case they exist; and (iii) analyzing the size of uniform interpolants. Clearly, these are fundamental steps on the way towards the computation and usage of uniform interpolation in practical applications. Regarding (i), we establish an intimate connection between uniform interpolants and the well-known notion of a bisimulation and characterize the existence of interpolants in terms of the existence of models with certain properties based on bisimulations. For (ii), our main result is that deciding the

existence of uniform interpolants is 2-EXPTIME-complete, and that methods for computing uniform interpolants on the level of concepts can be lifted to the TBox level. Finally, regarding (iii) we prove that the size of uniform interpolants is at most triple exponential in the size of the original TBox (upper bound), and that, in general, no shorter interpolants can be found (lower bound). In particular, this shows that the algorithm from [19] is flawed as it always yields uniform interpolants of at most double exponential size. Our methods, which are a mix of model-theory and automata-theory, also provide model-theoretic characterizations of conservative extensions, which are closely related to uniform interpolation [6]. Moreover, we use our approach to reprove the 2-EXPTIME upper bound for deciding conservative extensions from [6], in an alternative and arguably more transparent way.

Most proofs in this paper are deferred to the (appendix of the) long version, which is available at <http://www.csc.liv.ac.uk/~frank/publ/publ.html>

2 Getting Started

We introduce the description logic \mathcal{ALC} and define uniform interpolants and the dual notion of forgetting. Let N_C and N_R be disjoint and countably infinite sets of *concept* and *role names*. \mathcal{ALC} concepts are formed using the syntax rule

$$C, D \longrightarrow \top \mid A \mid \neg C \mid C \sqcap D \mid \exists r.C$$

where $A \in N_C$ and $r \in N_R$. The concept constructors \perp , \sqcup , and $\forall r.C$ are defined as abbreviations: \perp stands for $\neg\top$, $C \sqcup D$ for $\neg(\neg C \sqcap \neg D)$ and $\forall r.C$ abbreviates $\neg\exists r.\neg C$. A TBox is a finite set of *concept inclusions* $C \sqsubseteq D$, where C, D are \mathcal{ALC} -concepts. We use $C \equiv D$ as abbreviation for the two inclusions $C \sqsubseteq D$ and $D \sqsubseteq C$.

The semantics of \mathcal{ALC} -concepts is given in terms of *interpretations* $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, where $\Delta^{\mathcal{I}}$ is a non-empty set (the *domain*) and $\cdot^{\mathcal{I}}$ is the *interpretation function*, assigning to each $A \in N_C$ a set $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$, and to each $r \in N_R$ a relation $r^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$. The interpretation function is inductively extended to concepts as follows:

$$\begin{aligned} \top^{\mathcal{I}} &:= \Delta^{\mathcal{I}} & (\neg C)^{\mathcal{I}} &:= \Delta^{\mathcal{I}} \setminus C^{\mathcal{I}} & (C \sqcap D)^{\mathcal{I}} &:= C^{\mathcal{I}} \cap D^{\mathcal{I}} \\ (\exists r.C)^{\mathcal{I}} &:= \{d \in \Delta^{\mathcal{I}} \mid \exists e.(d, e) \in r^{\mathcal{I}} \wedge e \in C^{\mathcal{I}}\} \end{aligned}$$

An interpretation \mathcal{I} *satisfies* an inclusion $C \sqsubseteq D$ if $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$, and \mathcal{I} is a *model* of a TBox \mathcal{T} if it satisfies all inclusions in \mathcal{T} . A concept C is *subsumed* by a concept D relative to a TBox \mathcal{T} (written $\mathcal{T} \models C \sqsubseteq D$) if every model \mathcal{I} of \mathcal{T} satisfies the inclusion $C \sqsubseteq D$. We write $\mathcal{T} \models \mathcal{T}'$ to indicate that $\mathcal{T} \models C \sqsubseteq D$ for all $C \sqsubseteq D \in \mathcal{T}'$.

A set $\Sigma \subseteq N_C \cup N_R$ of concept and role names is called a *signature*. The signature $\text{sig}(C)$ of a concept C is the set of concept and role names occurring in C , and likewise for the signature $\text{sig}(C \sqsubseteq D)$ of an inclusion $C \sqsubseteq D$ and $\text{sig}(\mathcal{T})$ of a TBox \mathcal{T} . A Σ -TBox is a TBox with $\text{sig}(\mathcal{T}) \subseteq \Sigma$, and likewise for Σ -inclusions and Σ -concepts.

We now introduce the main notions studied in this paper: uniform interpolants and conservative extensions.

Definition 1. Let $\mathcal{T}, \mathcal{T}'$ be TBoxes and Σ a signature. \mathcal{T} and \mathcal{T}' are Σ -inseparable if for all Σ -inclusions $C \sqsubseteq D$, we have $\mathcal{T} \models C \sqsubseteq D$ iff $\mathcal{T}' \models C \sqsubseteq D$. We call

- \mathcal{T}' a *conservative extension* of \mathcal{T} if $\mathcal{T}' \supseteq \mathcal{T}$ and \mathcal{T} and \mathcal{T}' are Σ -inseparable for $\Sigma = \text{sig}(\mathcal{T})$.
- \mathcal{T} a *uniform Σ -interpolant* of \mathcal{T}' if $\text{sig}(\mathcal{T}) \subseteq \Sigma \subseteq \text{sig}(\mathcal{T}')$ and \mathcal{T} and \mathcal{T}' are Σ -inseparable.

Note that uniform Σ -interpolants are unique up to logical equivalence, if they exist.

The notion of forgetting as investigated in [19] is dual to uniform interpolation: a TBox \mathcal{T}' is *the result of forgetting about a signature Σ in a TBox \mathcal{T}* if \mathcal{T}' is a uniform $\text{sig}(\mathcal{T}) \setminus \Sigma$ -interpolant of \mathcal{T} .

Example 2. Let \mathcal{T} consist of the inclusions

- (1) $\text{Uni} \sqsubseteq \exists \text{has.st.Undergrad} \sqcap \exists \text{has.st.Grad}$
- (2) $\text{Uni} \sqcap \text{Undergrad} \sqsubseteq \perp$
- (3) $\text{Uni} \sqcap \text{Grad} \sqsubseteq \perp$
- (4) $\text{Undergrad} \sqcap \text{Grad} \sqsubseteq \perp$.

Then the TBox that consists of (2) and

$$\text{Uni} \sqsubseteq \exists \text{has.st.Undergrad} \sqcap \exists \text{has.st.}(\neg \text{Undergrad} \sqcap \neg \text{Uni})$$

is the result of forgetting $\{\text{Grad}\}$. Additionally forgetting Undergrad yields the TBox $\{\text{Uni} \sqsubseteq \exists \text{has.st.} \neg \text{Uni}\}$.

The following examples will be used to illustrate our characterizations. Proofs are provided once we have developed the appropriate tools.

Example 3. In the following, we always forget $\{B\}$.

- Let $\mathcal{T}_1 = \{A \sqsubseteq \exists r.B \sqcap \exists r.\neg B\}$ and $\Sigma_1 = \{A, r\}$. Then $\mathcal{T}'_1 = \{A \sqsubseteq \exists r.\top\}$ is a uniform Σ_1 -interpolant of \mathcal{T}_1 .
- Let $\mathcal{T}_2 = \{A \equiv B \sqcap \exists r.B\}$ and $\Sigma_2 = \{A, r\}$. Then $\mathcal{T}'_2 = \{A \sqsubseteq \exists r.(A \sqcup \neg \exists r.A)\}$ is a uniform Σ_2 -interpolant of \mathcal{T}_2 .
- For $\mathcal{T}_3 = \{A \sqsubseteq B, B \sqsubseteq \exists r.B\}$ and $\Sigma_3 = \{A, r\}$, there is no uniform Σ_3 -interpolant of \mathcal{T}_3 .
- For $\mathcal{T}_4 = \{A \sqsubseteq \exists r.B, A_0 \sqsubseteq \exists r.(A_1 \sqcap B), E \equiv A_1 \sqcap B \sqcap \exists r.(A_2 \sqcap B)\}$ and $\Sigma_4 = \{A, r, A_0, A_1, E\}$, there is no uniform Σ_4 -interpolant of \mathcal{T}_4 . Note that \mathcal{T}_4 is of a very simple form, namely an acyclic \mathcal{EL} -TBox, see [10].

Bisimulations are a central tool for studying the expressive power of \mathcal{ALC} , and play a crucial role also in our approach to uniform interpolants. We introduce them next. A *pointed interpretation* is a pair (\mathcal{I}, d) that consists of an interpretation \mathcal{I} and a $d \in \Delta^{\mathcal{I}}$.

Definition 4. Let Σ be a finite signature and $(\mathcal{I}_1, d_1), (\mathcal{I}_2, d_2)$ pointed interpretations. A relation $S \subseteq \Delta^{\mathcal{I}_1} \times \Delta^{\mathcal{I}_2}$ is a Σ -bisimulation between (\mathcal{I}_1, d_1) and (\mathcal{I}_2, d_2) if $(d_1, d_2) \in S$ and for all $(d, d') \in S$ the following conditions are satisfied:

1. $d \in A^{\mathcal{I}_1}$ iff $d' \in A^{\mathcal{I}_2}$, for all $A \in \Sigma \cap N_C$;
2. if $(d, e) \in r^{\mathcal{I}_1}$, then there exists $e' \in \Delta^{\mathcal{I}_2}$ such that $(d', e') \in r^{\mathcal{I}_2}$ and $(e, e') \in S$, for all $r \in \Sigma \cap N_R$;
3. if $(d', e') \in r^{\mathcal{I}_2}$, then there exists $e \in \Delta^{\mathcal{I}_1}$ such that $(d, e) \in r^{\mathcal{I}_1}$ and $(e, e') \in S$, for all $r \in \Sigma \cap N_C$.

(\mathcal{I}_1, d_1) and (\mathcal{I}_2, d_2) are Σ -bisimilar, written $(\mathcal{I}_1, d_1) \sim_{\Sigma} (\mathcal{I}_2, d_2)$, if there exists a Σ -bisimulation between them.

We now state the main connection between bisimulations and \mathcal{ALC} , well-known from modal logic [7]. Say that (\mathcal{I}_1, d_1) and (\mathcal{I}_2, d_2) are \mathcal{ALC}_Σ -equivalent, in symbols $(\mathcal{I}_1, d_1) \equiv_\Sigma (\mathcal{I}_2, d_2)$, if for all Σ -concepts C , $d_1 \in C^{\mathcal{I}_1}$ iff $d_2 \in C^{\mathcal{I}_2}$. An interpretation \mathcal{I} has *finite outdegree* if $\{d' \mid (d, d') \in \bigcup_{r \in \mathbb{N}_R} r^{\mathcal{I}}\}$ is finite, for all $d \in \Delta^{\mathcal{I}}$.

Theorem 5. *For all pointed interpretations (\mathcal{I}_1, d_1) and (\mathcal{I}_2, d_2) and all finite signatures Σ , $(\mathcal{I}_1, d_1) \sim_\Sigma (\mathcal{I}_2, d_2)$ implies $(\mathcal{I}_1, d_1) \equiv_\Sigma (\mathcal{I}_2, d_2)$; the converse holds for all $\mathcal{I}_1, \mathcal{I}_2$ of finite outdegree.*

Bisimulations enable a purely semantic characterization of uniform interpolants. For a pointed interpretation (\mathcal{I}, d) , we write $(\mathcal{I}, d) \models \exists_{\Sigma} \mathcal{T}$ when (\mathcal{I}, d) is Σ -bisimilar to some pointed interpretation (\mathcal{J}, d') with \mathcal{J} a model of \mathcal{T} . The notation reflects that what we express here can be understood as a form of bisimulation quantifier, see [5].

Theorem 6. *Let \mathcal{T} be a TBox and $\Sigma \subseteq \text{sig}(\mathcal{T})$. A Σ -TBox \mathcal{T}_Σ is a uniform Σ -interpolant of \mathcal{T} iff for all interpretations \mathcal{I} ,*

$$\mathcal{I} \models \mathcal{T}_\Sigma \iff \text{for all } d \in \Delta^{\mathcal{I}}, (\mathcal{I}, d) \models \exists_{\Sigma} \mathcal{T}. \quad (*)$$

For \mathcal{I} of finite outdegree, one can prove this result by employing compactness arguments and Theorem 5. To prove it in its full generality, we need the automata-theoretic machinery introduced in Section 4. We illustrate Theorem 6 by sketching a proof of Example 3(i). Correctness of 3(ii) is proved in the appendix, while 3(iii) and 3(iv) are addressed in Section 3. An interpretation \mathcal{I} is called a *tree interpretation* if the undirected graph $(\Delta^{\mathcal{I}}, \bigcup_{r \in \mathbb{N}_R} r^{\mathcal{I}})$ is a (possibly infinite) tree and $r^{\mathcal{I}} \cap s^{\mathcal{I}} = \emptyset$ for all distinct $r, s \in \mathbb{N}_R$.

Example 7. Let $\mathcal{T}_1 = \{A \sqsubseteq \exists r.B \sqcap \exists r.\neg B\}$, $\Sigma_1 = \{A, r\}$, and $\mathcal{T}'_1 = \{A \sqsubseteq \exists r.\top\}$ as in Example 3(i). We have $\mathcal{I} \models \mathcal{T}'_1$ iff $\forall d \in \Delta^{\mathcal{I}}: (\mathcal{I}, d) \sim_{\Sigma_1} (\mathcal{J}, d)$ for a tree model \mathcal{J} of \mathcal{T}'_1 iff $\forall d \in \Delta^{\mathcal{I}}: (\mathcal{I}, d) \sim_{\Sigma_1} (\mathcal{J}, d)$ for a tree interpretation \mathcal{J} such that $e \in A^{\mathcal{J}}$ implies $|\{d \mid (d, e) \in r^{\mathcal{J}}\}| \geq 2$ iff $\forall d \in \Delta^{\mathcal{I}}: (\mathcal{I}, d) \sim_{\Sigma_1} (\mathcal{J}, d)$ for a tree interpretation \mathcal{J} such that $e \in A^{\mathcal{J}}$ implies $e \in (\exists r.B \sqcap \exists r.\neg B)^{\mathcal{J}}$ iff $\forall d \in \Delta^{\mathcal{I}}: (\mathcal{I}, d) \models \exists_{\Sigma_1} \mathcal{T}_1$.

The first ‘iff’ relies on the fact that unraveling an interpretation into a tree interpretation preserves bisimilarity, the second one on the fact that bisimulations are oblivious to the duplication of successors, and the third one on a reinterpretation of $B \notin \Sigma_1$ in \mathcal{J} .

Theorem 6 also yields a characterization of conservative extensions in terms of bisimulations, which is as follows.

Theorem 8. *Let $\mathcal{T}, \mathcal{T}'$ be TBoxes. Then $\mathcal{T} \cup \mathcal{T}'$ is a conservative extension of \mathcal{T} iff for all interpretations \mathcal{I} , $\mathcal{I} \models \mathcal{T} \implies$ for all $d \in \Delta^{\mathcal{I}}$, $(\mathcal{I}, d) \models \exists_{\Sigma} \mathcal{T}'$ where $\Sigma = \text{sig}(\mathcal{T})$.*

3 Characterizing Existence of Interpolants

If we admit TBoxes that are infinite, then uniform Σ -interpolants always exist: for any TBox \mathcal{T} and signature Σ , the infinite TBox $\mathcal{T}_\Sigma^\infty$ that consists of all Σ -inclusions $C \sqsubseteq D$ with $\mathcal{T} \models C \sqsubseteq D$ is a uniform Σ -interpolant of \mathcal{T} . To refine this simple observation, we define the *role-depth* $\text{rd}(C)$ of a

concept C to be the nesting depth of existential restrictions in C . For every finite signature Σ and $m \geq 0$, one can fix a finite set $\mathcal{C}_f^m(\Sigma)$ of Σ -concepts D with $\text{rd}(D) \leq m$ such that every Σ -concept C with $\text{rd}(C) \leq m$ is equivalent to some $D \in \mathcal{C}_f^m(\Sigma)$. Let

$$\mathcal{T}_{\Sigma, m} = \{C \sqsubseteq D \mid \mathcal{T} \models C \sqsubseteq D \text{ and } C, D \in \mathcal{C}_f^m(\Sigma)\}.$$

Clearly, $\mathcal{T}_\Sigma^\infty$ is equivalent to $\bigcup_{m \geq 0} \mathcal{T}_{\Sigma, m}$ suggesting that if a uniform interpolant exists, it is one of the TBoxes $\mathcal{T}_{\Sigma, m}$. In fact, it is easy to see that the following are equivalent (this is similar to the approximation of uniform interpolants in [16]):

- (a) there does not exist a uniform Σ -interpolant of \mathcal{T} ;
- (b) no $\mathcal{T}_{\Sigma, m}$ is a uniform Σ -interpolant of \mathcal{T} ;
- (c) for all $m \geq 0$ there is a $k > m$ such that $\mathcal{T}_{\Sigma, m} \not\models \mathcal{T}_{\Sigma, k}$.

Our characterization of the (non)-existence of uniform interpolants is based on an analysis of the TBoxes $\mathcal{T}_{\Sigma, m}$. For an interpretation \mathcal{I} , $d \in \Delta^{\mathcal{I}}$, and $m \geq 0$, we use $\mathcal{I}^{\leq m}(d)$ to denote the *m-segment generated by d in \mathcal{I}* , i.e., the restriction of \mathcal{I} to those elements of $\Delta^{\mathcal{I}}$ that can be reached from d in at most m steps in the graph $(\Delta^{\mathcal{I}}, \bigcup_{r \in \mathbb{N}_R} r^{\mathcal{I}})$. Using the definition of $\mathcal{T}_{\Sigma, m}$, Theorem 5, and the fact that every m -segment can be described up to bisimulation using a concept of role-depth m , it can be shown that an interpretation \mathcal{I} is a model of $\mathcal{T}_{\Sigma, m}$ iff each of \mathcal{I} 's m -segments is Σ -bisimilar to an m -segment of a model of \mathcal{T} . Thus, if $\mathcal{T}_{\Sigma, m}$ is *not* a uniform interpolant, then this is due to a problem that cannot be ‘detected’ by m -segments, i.e., some Σ -part of a model of \mathcal{T} that is located before an m -segment can pose constraints on Σ -parts of the model after that segment, where ‘before’ and ‘after’ refer to reachability in $(\Delta^{\mathcal{I}}, \bigcup_{r \in \mathbb{N}_R} r^{\mathcal{I}})$.

The following result describes this in an exact way. Together with the equivalence of (a) and (b) above, it yields a first characterization of the existence of uniform interpolants. $\rho^{\mathcal{I}}$ denotes the root of a tree interpretation \mathcal{I} , $\mathcal{I}^{\leq m}$ abbreviates $\mathcal{I}^{\leq m}(\rho^{\mathcal{I}})$, and a Σ -tree interpretation is a tree interpretations that only interprets predicates from Σ .

Theorem 9. *Let \mathcal{T} be a TBox, $\Sigma \subseteq \text{sig}(\mathcal{T})$, and $m \geq 0$. Then $\mathcal{T}_{\Sigma, m}$ is not a uniform Σ -interpolant of \mathcal{T} iff*

($*_m$) *there exist two Σ -tree interpretations, \mathcal{I}_1 and \mathcal{I}_2 , of finite outdegree such that*

1. $\mathcal{I}_1^{\leq m} = \mathcal{I}_2^{\leq m}$;
2. $(\mathcal{I}_1, \rho^{\mathcal{I}_1}) \models \exists_{\Sigma} \mathcal{T}$;
3. $(\mathcal{I}_2, \rho^{\mathcal{I}_2}) \not\models \exists_{\Sigma} \mathcal{T}$;
4. For all successors d of $\rho^{\mathcal{I}_2}$: $(\mathcal{I}_2, d) \models \exists_{\Sigma} \mathcal{T}$.

Intuitively, Points 1 and 2 ensure that $\mathcal{I}_1^{\leq m} = \mathcal{I}_2^{\leq m}$ is an m -segment of a model of $\mathcal{T}_{\Sigma, m}$, Points 2 and 3 express that in models of \mathcal{T} , the Σ -part after the m -segment is constrained in some way, and Point 4 says that this is due to $\rho^{\mathcal{I}_1}$ and $\rho^{\mathcal{I}_2}$, i.e., the constraint is imposed ‘before’ the m -segment. The following example demonstrates how Theorem 9 can be used to prove non-existence of uniform interpolants.

Example 10. Let $\mathcal{T}_3 = \{A \sqsubseteq B, B \sqsubseteq \exists r.B\}$ and $\Sigma_3 = \{A, r\}$ as in Example 3(iii). We show that ($*_m$) holds for all

m and thus, there is no uniform Σ_3 -interpolant of \mathcal{T}_3 . Example 3(iv) is treated in the long version.

Let $m \geq 0$. Set $\mathcal{I}_1 = (\{0, 1, \dots\}, A^{\mathcal{I}_1}, r^{\mathcal{I}_1})$, where $A^{\mathcal{I}_1} = \{0\}$ and $r^{\mathcal{I}_1} = \{(n, n+1) \mid n \geq 0\}$, and let \mathcal{I}_2 be the restriction of \mathcal{I}_1 to $\{0, \dots, m\}$. Then (1) $\mathcal{I}_1^{\leq m} = \mathcal{I}_2^{\leq m}$; (2) $(\mathcal{I}_1, 0) \models \exists_{\Sigma_3}^{\sim} \mathcal{T}_3$ as the expansion of \mathcal{I}_1 by $B^{\mathcal{I}_1} = \{0, 1, \dots\}$ is a model of \mathcal{T}_3 ; (3) $(\mathcal{I}_2, 0) \not\models \exists_{\Sigma_3}^{\sim} \mathcal{T}_3$ as there is no infinite r -sequence in \mathcal{I}_2 starting at 0; and (4) $(\mathcal{I}_2, 1) \models \exists_{\Sigma_3}^{\sim} \mathcal{T}_3$ as the restriction of \mathcal{I}_2 to $\{1, \dots, m\}$ is a model of \mathcal{T}_3 .

The next example illustrates another use of Theorem 9 by identifying a class of signatures for which uniform interpolants always exist. Details are given in the long version.

Example 11 (Forgetting stratified concept names). A concept name A is *stratified* in \mathcal{T} if all occurrences of A in concepts from $\text{conc}(\mathcal{T}) = \{C, D \mid C \sqsubseteq D \in \mathcal{T}\}$ are exactly in nesting depth n of existential restrictions, for some $n \geq 0$. Let \mathcal{T} be a TBox and Σ a signature such that $\text{sig}(\mathcal{T}) \setminus \Sigma$ consists of stratified concept names only, i.e., we want to *forget* a set of stratified concept names. Then the existence of a uniform Σ -interpolant of \mathcal{T} is guaranteed; moreover, $\mathcal{T}_{\Sigma, m}$ is such an interpolant, where $m = \max\{\text{rd}(C) \mid C \in \text{conc}(\mathcal{T})\}$.

To turn Theorem 9 into a decision procedure for the existence of uniform interpolants, we prove that rather than testing $(*_m)$ for all m , it suffices to consider a single number m . This yields the final characterization of the existence of uniform interpolants. We use $|\mathcal{T}|$ to denote the *length* of a TBox \mathcal{T} , i.e., the number of symbols needed to write it.

Theorem 12. *Let \mathcal{T} be a TBox and $\Sigma \subseteq \text{sig}(\mathcal{T})$. Then there does not exist a uniform Σ -interpolant of \mathcal{T} iff $(*_M_{\mathcal{T}+1})$ from Theorem 9 holds, where $M_{\mathcal{T}} := 2^{2^{|\mathcal{T}|}}$.*

It suffices to show that $(*_M_{\mathcal{T}+1})$ implies $(*_m)$ for all $m \geq M_{\mathcal{T}} + 1$. The proof idea is as follows. Denote by $\text{cl}(\mathcal{T})$ the closure under single negation and subconcepts of $\text{conc}(\mathcal{T})$. The type of some $d \in \Delta^{\mathcal{I}}$ in an interpretation \mathcal{I} is

$$\text{tp}^{\mathcal{I}}(d) := \{C \in \text{cl}(\mathcal{T}) \mid d \in C^{\mathcal{I}}\}.$$

Many constructions for \mathcal{ALC} (such as blocking in tableaux, filtrations of interpretations, etc.) exploit the fact that the relevant information about any element d in an interpretation is given by its type. This can be exploited e.g. to prove EXPTIME upper bounds as there are ‘only’ exponentially many distinct types. In the proof of Theorem 12, we make use of a ‘pumping lemma’ that enables us to transform any pair $\mathcal{I}_1, \mathcal{I}_2$ witnessing $(*_M_{\mathcal{T}+1})$ into a witness $\mathcal{I}'_1, \mathcal{I}'_2$ for $(*_m)$ when $m \geq M_{\mathcal{T}} + 1$. The construction depends on the relevant information about elements of $\Delta^{\mathcal{I}_1}$ and $\Delta^{\mathcal{I}_2}$; in contrast to standard constructions, however, types are not sufficient and must be replaced by *extension sets* $\text{Ext}^{\mathcal{I}}(d)$, defined as

$$\text{Ext}^{\mathcal{I}}(d) = \{\text{tp}^{\mathcal{I}}(d') \mid \exists \mathcal{J} : \mathcal{J} \models \mathcal{T} \text{ and } (\mathcal{I}, d) \sim_{\Sigma} (\mathcal{J}, d')\}$$

and capturing all ways in which the restriction of $\text{tp}^{\mathcal{I}}(d)$ to Σ -concepts can be extended to a full type in models of \mathcal{T} . As the number of such extension sets is double exponential in $|\mathcal{T}|$ and we have to consider pairs $(d_1, d_2) \in \Delta^{\mathcal{I}_1} \times \Delta^{\mathcal{I}_2}$, we (roughly) obtain a $M_{\mathcal{T}}^2$ bound. Details are in the long version.

We note that, by Theorem 12, the uniform Σ -interpolant of a TBox \mathcal{T} exists iff $\mathcal{T} \cup \mathcal{T}_{\Sigma, M_{\mathcal{T}}^2+1}$ is a conservative extension of $\mathcal{T}_{\Sigma, M_{\mathcal{T}}^2+1}$. With the decidability of conservative extensions proved in [6], this yields decidability of the existence of uniform interpolants. However, the size of $\mathcal{T}_{\Sigma, M_{\mathcal{T}}^2+1}$ is non-elementary, and so is the running time of the resulting algorithm. We next show how to improve this.

4 Automata Constructions / Complexity

We develop a worst-case optimal algorithm for deciding the existence of uniform interpolants in \mathcal{ALC} , exploiting Theorem 12 and making use of alternating automata. As a by-product, we prove the fundamental characterization of uniform interpolants in terms of bisimulation stated as Theorem 6 without the initial restriction to interpretations of finite outdegree. We also obtain a representation of uniform interpolants as automata and a novel, more transparent proof of the 2-EXPTIME upper bound for deciding conservative extensions originally established in [6].

We use amorphous alternating parity tree automata in the style of Wilke [20], which run on unrestricted interpretations rather than on trees, only. We call them *tree* automata as they are in the tradition of more classical forms of such automata. In particular, a run of an automaton is tree-shaped, even if the input interpretation is not.

Definition 13 (APTA). An *alternating parity tree automaton* (APTA) is a tuple $\mathcal{A} = (Q, \Sigma_N, \Sigma_E, q_0, \delta, \Omega)$, where Q is a finite set of states, $\Sigma_N \subseteq \mathbb{N}_C$ is the finite node alphabet, $\Sigma_E \subseteq \mathbb{N}_R$ is the finite edge alphabet, $q_0 \in Q$ is the initial state, $\delta : Q \rightarrow \text{mov}(\mathcal{A})$, is the transition function with $\text{mov}(\mathcal{A}) = \{\text{true}, \text{false}, A, \neg A, q, q \wedge q', q \vee q', \langle r \rangle q, [r]q \mid A \in \Sigma_N, q, q' \in Q, r \in \Sigma_E\}$ the set of moves of the automaton, and $\Omega : Q \rightarrow \mathbb{N}$ is the priority function.

Intuitively, the move q means that the automaton sends a copy of itself in state q to the element of the interpretation that it is currently processing, $\langle r \rangle q$ means that a copy in state q is sent to an r -successor of the current element, and $[r]q$ means that a copy in state q is sent to every r -successor.

It will be convenient to use arbitrary modal formulas in negation normal form when specifying the transition function of APTAs. The more restricted form required by Definition 13 can then be attained by introducing intermediate states. In subsequent constructions that involve APTAs, we will not describe those additional states explicitly. However, we will (silently) take them into account when stating size bounds for automata.

In what follows, a Σ -labelled tree is a pair (T, ℓ) with T a tree and $\ell : T \rightarrow \Sigma$ a node labelling function. A path π in a tree T is a subset of T such that $\varepsilon \in \pi$ and for each $x \in \pi$ that is not a leaf in T , π contains one son of x .

Definition 14 (Run). Let (\mathcal{I}, d_0) be a pointed $\Sigma_N \cup \Sigma_E$ -interpretation and $\mathcal{A} = (Q, \Sigma_N, \Sigma_E, q_0, \delta, \Omega)$ an APTA. A run of \mathcal{A} on (\mathcal{I}, d_0) is a $Q \times \Delta^{\mathcal{I}}$ -labelled tree (T, ℓ) such that $\ell(\varepsilon) = (q_0, d_0)$ and for every $x \in T$ with $\ell(x) = (q, d)$:

- $\delta(q) \neq \text{false}$;
- if $\delta(q) = A (\delta(q) = \neg A)$, then $d \in A^{\mathcal{I}} (d \notin A^{\mathcal{I}})$;

- if $\delta(q) = q' \wedge q''$, then there are sons y, y' of x with $\ell(y) = (q', d)$ and $\ell(y') = (q'', d)$;
- if $\delta(q) = q' \vee q''$, then there is a son y of x with $\ell(y) = (q', d)$ or $\ell(y) = (q'', d)$;
- if $\delta(q) = \langle r \rangle q'$, then there is a $(d, d') \in r^{\mathcal{I}}$ and a son y of x with $\ell(y) = (q', d')$;
- if $\delta(q) = [r]q'$ and $(d, d') \in r^{\mathcal{I}}$, then there is a son y of x with $\ell(y) = (q', d')$.

A run (T, ℓ) is *accepting* if for every path π of T , the maximal $i \in \mathbb{N}$ with $\{x \in \pi \mid \ell(x) = (q, d) \text{ with } \Omega(q) = i\}$ infinite is even. We use $L(\mathcal{A})$ to denote the language accepted by \mathcal{A} , i.e., the set of pointed $\Sigma_N \cup \Sigma_E$ -interpretations (\mathcal{I}, d) such that there is an accepting run of \mathcal{A} on (\mathcal{I}, d) .

Using the fact that runs are always tree-shaped, it is easy to prove that the languages accepted by APTAs are closed under $\Sigma_N \cup \Sigma_E$ -bisimulations. It is this property that makes this automaton model particularly useful for our purposes. APTAs can be complemented in polytime in the same way as other alternating tree automata, and for all APTAs \mathcal{A}_1 and \mathcal{A}_2 , one can construct in polytime an APTA that accepts $L(\mathcal{A}_1) \cap L(\mathcal{A}_2)$. Wilke shows that the emptiness problem for APTAs is EXPTIME-complete [20].

We now show that uniform Σ -interpolants of \mathcal{ALC} -TBoxes can be *represented as APTAs*, in the sense of the following theorem and of Theorem 6.

Theorem 15. *Let \mathcal{T} be a TBox and $\Sigma \subseteq \text{sig}(\mathcal{T})$ a signature. Then there exists an APTA $\mathcal{A}_{\mathcal{T}, \Sigma} = (Q, \Sigma \cap \mathbf{N}_C, \Sigma \cap \mathbf{N}_R, q_0, \delta, \Omega)$ with $|Q| \in 2^{\mathcal{O}(|\mathcal{T}|)}$ such that $L(\mathcal{A}_{\mathcal{T}, \Sigma})$ consists of all pointed Σ -interpretations (\mathcal{I}, d) with $(\mathcal{I}, d) \models \exists_{\Sigma}^{\sim} \mathcal{T}$. $\mathcal{A}_{\mathcal{T}, \Sigma}$ can be constructed in time $2^{p(|\mathcal{T}|)}$, p a polynomial.*

The construction of the automaton $\mathcal{A}_{\mathcal{T}, \Sigma}$ from Theorem 15 resembles the construction of uniform interpolants in the μ -calculus using non-deterministic automata described in [2], but is transferred to TBoxes and alternating automata.

Fix a TBox \mathcal{T} and a signature Σ and assume w.l.o.g. that \mathcal{T} has the form $\{\top \sqsubseteq C_{\mathcal{T}}\}$, with $C_{\mathcal{T}}$ in negation normal form [1]. Recall the notion of a type introduced in Section 3. We use $\text{TP}(\mathcal{T})$ to denote the set of all types realized in some model of \mathcal{T} , i.e., $\text{TP}(\mathcal{T}) = \{\text{tp}^{\mathcal{I}}(d) \mid \mathcal{I} \text{ model of } \mathcal{T}, d \in \Delta^{\mathcal{I}}\}$. Note that $\text{TP}(\mathcal{T})$ can be computed in time exponential in the size of \mathcal{T} since concept satisfiability w.r.t. TBoxes is EXPTIME-complete in \mathcal{ALC} [1]. Given $t, t' \in \text{TP}(\mathcal{T})$ and $r \in \Sigma$, we write $t \rightsquigarrow_r t'$ if $C \in t'$ implies $\exists r.C \in t$ for all $\exists r.C \in \text{cl}(\mathcal{T})$. Now define the automaton $\mathcal{A}_{\mathcal{T}, \Sigma} := (Q, \Sigma_N, \Sigma_E, q_0, \delta, \Omega)$, where

$$\begin{aligned} Q &= \text{TP}(\mathcal{T}) \uplus \{q_0\} & \Sigma_N &= \Sigma \cap \mathbf{N}_C & \Sigma_E &= \Sigma \cap \mathbf{N}_R \\ \delta(q_0) &= \bigvee \text{TP}(\mathcal{T}) \\ \delta(t) &= \bigwedge_{A \in t \cap \mathbf{N}_C \cap \Sigma} A \wedge \bigwedge_{A \in (\mathbf{N}_C \cap \Sigma) \setminus t} \neg A \\ &\wedge \bigwedge_{r \in \Sigma \cap \mathbf{N}_R} [r] \bigvee \{t' \in \text{TP}(\mathcal{T}) \mid t \rightsquigarrow_r t'\} \\ &\wedge \bigwedge_{\exists r.C \in t, r \in \Sigma} \langle r \rangle \bigvee \{t' \in \text{TP}(\mathcal{T}) \mid C \in t \wedge t \rightsquigarrow_r t'\} \\ \Omega(q) &= 0 \text{ for all } q \in Q \end{aligned}$$

Here, the empty conjunction represents true and the empty disjunction represents false. The acceptance condition of the automaton is trivial, which (potentially) changes when we complement it subsequently. We prove in the appendix that this automaton satisfies the conditions in Theorem 15.

We now develop a decision procedure for the existence of uniform interpolants by showing that the characterization of the existence of uniform interpolants provided by Theorem 12 can be captured by APTAs, in the following sense.

Theorem 16. *Let \mathcal{T} be a TBox, $\Sigma \subseteq \text{sig}(\mathcal{T})$ a signature, and $m \geq 0$. Then there is an APTA $\mathcal{A}_{\mathcal{T}, \Sigma, m} = (Q, \Sigma_N, \Sigma_E, q_0, \delta, \Omega)$ such that $L(\mathcal{A}) \neq \emptyset$ iff Condition $(*_m)$ from Theorem 9 is satisfied. Moreover, $|Q| \in \mathcal{O}(2^{\mathcal{O}(n)} + \log^2 m)$ and $|\Sigma_N|, |\Sigma_E| \in \mathcal{O}(n + \log m)$, where $n = |\mathcal{T}|$.*

The size of $\mathcal{A}_{\mathcal{T}, \Sigma, m}$ is exponential in $|\mathcal{T}|$ and logarithmic in m . By Theorem 12, we can set $m = 2^{2^{|\mathcal{T}|}}$, and thus the size of $\mathcal{A}_{\mathcal{T}, \Sigma, m}$ is exponential in $|\mathcal{T}|$. Together with the EXPTIME emptiness test for APTAs, we obtain a 2-EXPTIME decision procedure for the existence of uniform interpolants. We construct $\mathcal{A}_{\mathcal{T}, \Sigma, m}$ as an intersection of four APTAs, each ensuring one of the conditions of $(*_m)$; building the automaton for Condition 2 involves complementation. The automaton $\mathcal{A}_{\mathcal{T}, \Sigma, m}$ runs over an extended alphabet that allows to encode both of the interpretations \mathcal{I}_1 and \mathcal{I}_2 mentioned in $(*_m)$, plus a ‘depth counter’ for enforcing Condition 1 of $(*_m)$.

A similar, but simpler construction can be used to reprove the 2-EXPTIME upper bound for deciding conservative extensions established in [6]. The construction only depends on Theorem 8, but not on the material in Section 3 and is arguably more transparent than the original one.

Theorem 17. *Given TBoxes \mathcal{T} and \mathcal{T}' , it can be decided in time $2^{p(|\mathcal{T}| \cdot 2^{|\mathcal{T}'|})}$ whether $\mathcal{T} \cup \mathcal{T}'$ is a conservative extension of \mathcal{T} , for some polynomial $p()$.*

A 2-EXPTIME lower bound was also established in [6], thus the upper bound stated in Theorem 17 is tight. This lower bound transfers to the existence of uniform interpolants: one can show that $\mathcal{T}' = \mathcal{T} \cup \{\top \sqsubseteq C\}$ is a conservative extension of \mathcal{T} iff there is a uniform $\text{sig}(\mathcal{T}) \cup \{r\}$ -interpolant of $\mathcal{T}_0 = \mathcal{T} \cup \{\neg C \sqsubseteq A, A \sqsubseteq \exists r.A\} \cup \{\exists s.A \sqsubseteq A \mid s \in \text{sig}(\mathcal{T}')\}$, with r, A are fresh. This yields the main result of this section.

Theorem 18. *It is 2-EXPTIME-complete to decide, given a TBox \mathcal{T} and a signature $\Sigma \subseteq \text{sig}(\mathcal{T})$, whether there exists a uniform Σ -interpolant of \mathcal{T} .*

5 Computing Interpolants / Interpolant Size

We show how to compute smaller uniform interpolants than the non-elementary $\mathcal{T}_{\Sigma, M_{\mathcal{T}}^2 + 1}$ and establish a matching upper bound on their size. Let C be a concept and $\Sigma \subseteq \text{sig}(C)$ a signature. A concept C' is called a *concept uniform Σ -interpolant of C* if $\text{sig}(C) \subseteq \Sigma$, $\emptyset \models C \sqsubseteq C'$, and $\emptyset \models C' \sqsubseteq D$ for every concept D such that $\text{sig}(D) \subseteq \Sigma$ and $\emptyset \models C \sqsubseteq D$. The following result is proved in [14].

Theorem 19. *For every concept C and signature $\Sigma \subseteq \text{sig}(C)$ one can effectively compute a concept uniform Σ -interpolant C' of C of at most exponential size in C .*

This result can be lifted to (TBox) uniform interpolants by ‘internalization’ of the TBox. This is very similar to what is attempted in [19], but we use different bounds on the role depth of the internalization concepts. More specifically, let $\mathcal{T} = \{\top \sqsubseteq C_{\mathcal{T}}\}$ have a uniform Σ -interpolant and R denote the set of role names in \mathcal{T} . For a concept C , define inductively

$$\forall R^{\leq 0}.C = C, \quad \forall R^{\leq n+1}.C = C \sqcap \prod_{r \in R} \forall r. \forall R^{\leq n}.C$$

It can be shown using Theorem 12 that for $m = 2^{2^{|\mathcal{T}|+1}} + 2^{|\mathcal{T}|} + 2$ and C a concept uniform Σ -interpolant of $\forall R^{\leq m}.C_{\mathcal{T}}$, the TBox $\mathcal{T}' = \{\top \sqsubseteq C\}$ is a uniform Σ -interpolant of \mathcal{T} . A close inspection of the construction underlying the proof of Theorem 19 applied to $\forall R^{\leq m}.C_{\mathcal{T}}$ reveals that $\text{rd}(C) \leq \text{rd}(\forall R^{\leq m}.C_{\mathcal{T}})$ and that the size of C is at most triple exponential in $|\mathcal{T}|$. This yields the following upper bound.

Theorem 20. *Let \mathcal{T} be an \mathcal{ALC} -TBox and $\Sigma \subseteq \text{sig}(\mathcal{T})$. If there is a uniform Σ -interpolant of \mathcal{T} , then there is one of size at most $2^{2^{2^{p(|\mathcal{T}|)}}$, p a polynomial.*

A matching lower bound on the size of uniform interpolants can be obtained by transferring a lower bound on the size of so-called witness concepts for (non-)conservative extensions established in [6]:

Theorem 21. *There exists a signature Σ of cardinality 4 and a family of TBoxes $(\mathcal{T}_n)_{n>0}$ such that, for all $n > 0$,*

- (i) $|\mathcal{T}_n| \in \mathcal{O}(n^2)$ and
- (ii) every uniform Σ -interpolant $\{\top \sqsubseteq C_{\mathcal{T}}\}$ for \mathcal{T}_n is of size at least $2^{(2^n \cdot 2^{2^n}) - 2}$.

6 Conclusions

We view the characterizations, tools, and results obtained in this paper as a general foundation for working with uniform interpolants in expressive DLs. In fact, we believe that the established framework can be extended to other expressive DLs such as \mathcal{ALC} extended with number restrictions and/or inverse roles without too many hassles: the main modifications required should be a suitable modification of the notion of bisimulation and (at least in the case of number restrictions) a corresponding extension of the automata model. Other extensions, such as with nominals, require more efforts.

In concrete applications, what to do when the desired uniform Σ -interpolant does not exist? In applications such as *ontology re-used* and *ontology summary*, one option is to extend the signature Σ , preferably in a minimal way, and then to use the interpolant for the extended signature. We believe that Theorem 9 can be helpful to investigate this further, loosely in the spirit of Example 11. In applications such as *predicate hiding*, an extension of Σ might not be acceptable. It is then possible to resort to a more expressive DL in which uniform interpolants always exist. In fact, Theorem 15 and the fact that APTAs have the same expressive power as the μ -calculus [20] point the way towards the extension of \mathcal{ALC} with fixpoint operators.

References

- [1] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. Patel-Schneider. *The Description Logic Handbook*. Cambridge University Press, 2003.
- [2] G. D’Agostino and M. Hollenberg. Uniform interpolation, automata and the modal μ -calculus. In *Advances in Modal Logic, Volume 1*, pages 73–84. CSLI Publications, 1998.
- [3] D.E. Muller and P.E. Schupp. Alternating automata on infinite trees. *Theoretical Computer Science*, 54:267–276, 1987.
- [4] T. Eiter, G. Ianni, R. Schindlauer, H. Tompits, and K. Wang. Forgetting in managing rules and ontologies. In *Web Intelligence*, pages 411–419, 2006.
- [5] T. French. *Bisimulation quantifiers for modal logics*. PhD thesis, University of Western Australia, 2006.
- [6] S. Ghilardi, C. Lutz, and F. Wolter. Did I damage my ontology? A case for conservative extensions in description logic. In *Proc. of KR06*, pages 187–197. AAAI Press, 2006.
- [7] V. Goranko and M. Otto. Model theory of modal logic. In *Handbook of Modal Logic*. Elsevier, 2007.
- [8] B. Cuenca Grau and B. Motik. Pushing the limits of reasoning over ontologies with hidden content. In *Proc. of KR10*, AAAI Press, 2010.
- [9] A. Herzig and J. Mengin. Uniform interpolation by resolution in modal logic. In *Proc. of JELIA08*, volume 5293 of LNCS, pages 219–231. Springer, 2008.
- [10] B. Konev, D. Walther, and F. Wolter. Forgetting and uniform interpolation in large-scale description logic terminologies. In *Proc. of IJCAI09*, pages 830–835, 2009.
- [11] R. Kontchakov, F. Wolter, and M. Zakharyashev. Logic-based ontology comparison and module extraction, with an application to DL-Lite. *Artificial Intelligence*, 174(15):1093–1141, 2010.
- [12] J. Lang, P. Liberatore, and P. Marquis. Propositional independence: formula-variable independence and forgetting. *J. of Artificial Intelligence Research*, 18:391–443, 2003.
- [13] R. Reiter and F. Lin. Forget it! In *Proc. of AAAI Fall Symposium on Relevance*, pages 154–159, 1994.
- [14] B. ten Cate, W. Conradie, M. Marx, and Y. Venema. Definitorially complete description logics. In *Proc. of KR06*, pages 79–89. AAAI Press, 2006.
- [15] A. Visser. Uniform interpolation and layered bisimulation. In *Proc. of Gödel96*, volume 6 of *Lecture notes in logic*, pages 139–164. Springer Verlag, 1996.
- [16] K. Wang, Z. Wang, R.W. Topor, J.Z. Pan, and G. Antoniou. Concept and role forgetting in \mathcal{ALC} ontologies. In *Proc. of ISWC09*, volume 5823 of LNCS, pages 666–681. Springer, 2009.
- [17] Z. Wang, K. Wang, R. W. Topor, and J. Z. Pan. Forgetting concepts in DL-Lite. In *Proc. of ESWC08*, volume 5021 of LNCS, pages 245–257. Springer, 2008.
- [18] Z. Wang, K. Wang, R.W. Topor, J.Z. Pan, and G. Antoniou. Uniform interpolation for \mathcal{ALC} revisited. In *Proc. of ACAI09*, volume 5866 of LNCS, pages 528–537. Springer, 2009.
- [19] Z. Wang, K. Wang, R. W. Topor, and X. Zhang. Tableau-based forgetting in \mathcal{ALC} ontologies. In *Proc. of ECAI10*, pages 47–52, IOS Press, 2010.
- [20] T. Wilke. Alternating tree automata, parity games, and modal μ -calculus. *Bulletin of the Belg. Math. Soc.*, 8(2), 2001.

7 Proofs for Section 2

Proof sketch for Example 3 (ii) Recall that $\mathcal{T}_2 = \{A \equiv B \sqcap \exists r.B\}$ and $\Sigma_2 = \{A, r\}$. We show that $\mathcal{T}'_2 = \{A \sqsubseteq \exists r.(A \sqcup \neg \exists r.A)\}$ is a uniform Σ_2 -interpolant of \mathcal{T}_2 . To this end, we prove the criterion of Theorem 6.

Let \mathcal{I} be a model of $A \sqsubseteq \exists r.(A \sqcup \neg \exists r.A)$. Then every (\mathcal{I}, d) , $d \in \Delta^{\mathcal{I}}$, is bisimilar to a tree-interpretation $(\mathcal{J}, \rho^{\mathcal{J}})$ that is a model of $A \sqsubseteq \exists r.(A \sqcup \neg \exists r.A)$. We define a new interpretation \mathcal{J}' that coincides with \mathcal{J} except that B is interpreted as follows: for every $e \in \Delta^{\mathcal{J}}$ with $e \in A^{\mathcal{J}}$ let $e \in B^{\mathcal{J}'}$ and, if there does not exist an r -successor of e in $A^{\mathcal{J}}$, then take an r -successor e' of e with $e' \notin (\exists r.A)^{\mathcal{J}}$ and let $e' \in B^{\mathcal{J}'}$ as well. Such an e' exists since \mathcal{J} is a model of $A \sqsubseteq \exists r.(A \sqcup \neg \exists r.A)$. It is readily checked that \mathcal{J}' is a model of \mathcal{T}_2 . Thus, for all $d \in \Delta^{\mathcal{I}}$, $(\mathcal{I}, d) \models \exists_{\Sigma_2}^{\sim} \mathcal{T}_2$.

Conversely, let $(\mathcal{I}, d) \models \exists_{\Sigma_2}^{\sim} \mathcal{T}_2$, for all $d \in \Delta^{\mathcal{I}}$. Let $d \in A^{\mathcal{I}}$ and assume $d \notin \exists r.(A \sqcup \neg \exists r.A)$. Then, no r -successor of d is in $A^{\mathcal{I}}$ and all r -successors of d have an r -successor in $A^{\mathcal{I}}$. Let $(\mathcal{I}, d) \sim_{\Sigma_2} (\mathcal{J}, d')$ with \mathcal{J} a model of \mathcal{T}_2 . Then $d' \in A^{\mathcal{J}}$, no r -successor of d' is in $A^{\mathcal{J}}$, and all r -successors of d' have an r -successor in $A^{\mathcal{J}}$. We have $d' \in (B \sqcap \exists r.B)^{\mathcal{J}}$ and so $d' \in B^{\mathcal{J}}$ and there exists an r -successor d'' of d' such that $d'' \in B^{\mathcal{J}}$. Since $d'' \notin A^{\mathcal{J}}$, there does not exist an r -successor of d'' that is in $B^{\mathcal{J}}$. But then no r -successor of d'' is in $A^{\mathcal{J}}$ and we have derived a contradiction.

To prove Theorem 6 in its full generality, we will rely on the automata-theoretic machinery introduced in Section 4. For now, we only establish a modified version where “for all interpretation \mathcal{I} ” is replaced with “for all interpretations \mathcal{I} of finite outdegree”.

Theorem 6 (Modified Version) Let \mathcal{T} be a TBox and $\Sigma \subseteq \text{sig}(\mathcal{T})$. A Σ -TBox \mathcal{T}_{Σ} is a uniform Σ -interpolant of \mathcal{T} iff for all interpretations \mathcal{I} of finite outdegree,

$$\mathcal{I} \models \mathcal{T}_{\Sigma} \Leftrightarrow \text{for all } d \in \Delta^{\mathcal{I}}, (\mathcal{I}, d) \models \exists_{\Sigma}^{\sim} \mathcal{T}. \quad (*)$$

Proof. “if”. Assume that (*) is satisfied for all interpretations \mathcal{I} of finite outdegree. We have to show that for all Σ -inclusions $C \sqsubseteq D$,

$$\mathcal{T}_{\Sigma} \models C \sqsubseteq D \text{ iff } \mathcal{T} \models C \sqsubseteq D.$$

Let $\text{sig}(C \sqsubseteq D) \subseteq \Sigma$ and assume $\mathcal{T}_{\Sigma} \not\models C \sqsubseteq D$. Then $C \sqcap \neg D$ is satisfiable w.r.t. \mathcal{T}_{Σ} , i.e., there is a pointed model (\mathcal{I}, d) of \mathcal{T}_{Σ} with finite outdegree and $d \in (C \sqcap \neg D)^{\mathcal{I}}$. By (*), there is thus a pointed model (\mathcal{J}, e) of \mathcal{T} with $(\mathcal{I}, d) \sim_{\Sigma} (\mathcal{J}, e)$. Together with $d \in (C \sqcap \neg D)^{\mathcal{I}}$, the latter implies $e \in (C \sqcap \neg D)^{\mathcal{J}}$. Thus, $C \sqcap \neg D$ is satisfiable w.r.t. \mathcal{T} , implying $\mathcal{T} \not\models C \sqsubseteq D$. Conversely, let $\mathcal{T} \not\models C \sqsubseteq D$. Then there is a pointed model (\mathcal{I}, d) of \mathcal{T} with finite outdegree and $d \in (C \sqcap \neg D)^{\mathcal{I}}$. Trivially, \mathcal{I} satisfies the right-hand side of (*), whence $\mathcal{I} \models \mathcal{T}_{\Sigma}$ and we are done.

For the “only if” direction, we first need a preliminary. An interpretation \mathcal{I} is *modally saturated* iff it satisfies the following condition, for all $r \in \mathbb{N}_{\mathbb{R}}$: if $d \in \Delta^{\mathcal{I}}$ and Γ is a (potentially infinite) set of concepts such that, for all finite $\Psi \subseteq \Gamma$, there is a d' with $(d, d') \in r^{\mathcal{I}}$ and $d' \in \Psi^{\mathcal{I}}$, then there is an e with $(d, e) \in r^{\mathcal{I}}$ and $e \in \Gamma^{\mathcal{I}}$. The most important facts about modally saturated interpretations we need

here are (i) every (finite) or infinite set of concepts that is satisfiable w.r.t. a TBox \mathcal{T} is satisfiable in a modally saturated model of \mathcal{T} ; (ii) every interpretation with finite outdegree is modally saturated; and (iii) Point 2 of Theorem 5 can be generalized from interpretations of finite outdegree to modally saturated interpretations [7].

“only if”. Assume that \mathcal{T}_{Σ} is a uniform Σ -interpolant of \mathcal{T} . First assume that \mathcal{I} is an interpretation that satisfies the right-hand side of (*). By Point 1 of Theorem 5, $\mathcal{I} \models \mathcal{T}$ and since $\mathcal{T} \models \mathcal{T}_{\Sigma}$, also $\mathcal{I} \models \mathcal{T}_{\Sigma}$. Now assume that \mathcal{I} is a model of \mathcal{T}_{Σ} of finite outdegree and let $d \in \Delta^{\mathcal{I}}$. Define Γ to be the set of all Σ -concepts C with $d \in C^{\mathcal{I}}$. Clearly, every finite subset $\Gamma' \subseteq \Gamma$ is satisfiable w.r.t. \mathcal{T}_{Σ} . Since \mathcal{T}_{Σ} is a uniform Σ -interpolant of \mathcal{T} , every such Γ' is also satisfiable w.r.t. \mathcal{T} (since $\mathcal{T}_{\Sigma} \not\models \top \sqsubseteq \neg \bigwedge \Gamma'$ implies $\mathcal{T} \not\models \top \sqsubseteq \neg \bigwedge \Gamma'$). By compactness of \mathcal{ALC} , Γ is satisfiable w.r.t. \mathcal{T} . By (i), there thus exists a modally saturated pointed model (\mathcal{J}, e) of \mathcal{T} such that $e \in \Gamma^{\mathcal{J}}$. Since $d \in \Gamma^{\mathcal{I}}$ and $e \in \Gamma^{\mathcal{J}}$ (and since, by definition, Γ contains each Σ -concept or its negation), we have $d \in C^{\mathcal{I}}$ iff $e \in C^{\mathcal{J}}$ for all \mathcal{ALC} -concepts C over Σ . By (ii) and (iii), this yields $\mathcal{I} \sim_{\Sigma} \mathcal{J}$ and we are done. \square \square

8 Proofs for Section 3

For a tree interpretation \mathcal{I} and $d \in \Delta^{\mathcal{I}}$ we denote by $\mathcal{I}(d)$ the tree interpretation induced by the subtree generated by d in \mathcal{I} .

Besides of \mathcal{ALC}_{Σ} -equivalence, we now also require a characterization of \mathcal{ALC}_{Σ} -equivalence for concepts of roles depth bounded by some m .

Two pointed interpretations are *\mathcal{ALC}_{Σ} - m -equivalent*, in symbols $(\mathcal{I}_1, d_1) \equiv_{\Sigma}^m (\mathcal{I}_2, d_2)$, if, and only if, for all Σ -concepts C with $\text{rd}(C) \leq m$, $d_1 \in C^{\mathcal{I}_1}$ iff $d_2 \in C^{\mathcal{I}_2}$.

The corresponding model-theoretic notion is that of *m -bisimilarity*, which is defined inductively as follows: (\mathcal{I}_1, d_1) and (\mathcal{I}_2, d_2) are

- $(\Sigma, 0)$ -bisimilar, in symbols $(\mathcal{I}_1, d_1) \sim_{\Sigma}^0 (\mathcal{I}_2, d_2)$, if $d_1 \in A^{\mathcal{I}_1}$ iff $d_2 \in A^{\mathcal{I}_2}$ for all $A \in \Sigma \cap \mathbb{N}_{\mathbb{C}}$.
- $(\Sigma, n+1)$ -bisimilar, in symbols $(\mathcal{I}_1, d_1) \sim_{\Sigma}^{n+1} (\mathcal{I}_2, d_2)$, if $(\mathcal{I}_1, d_1) \sim_{\Sigma}^0 (\mathcal{I}_2, d_2)$ and
 - for all $(d_1, e_1) \in r^{\mathcal{I}_1}$ there exists $e_2 \in \Delta^{\mathcal{I}_2}$ such that $(d_2, e_2) \in r^{\mathcal{I}_2}$ and $(\mathcal{I}_1, e_1) \sim_{\Sigma}^n (\mathcal{I}_2, e_2)$, for all $r \in \Sigma$;
 - for all $(d_2, e_2) \in r^{\mathcal{I}_2}$ there exists $e_1 \in \Delta^{\mathcal{I}_1}$ such that $(d_1, e_1) \in r^{\mathcal{I}_1}$ and $(\mathcal{I}_1, e_1) \sim_{\Sigma}^n (\mathcal{I}_2, e_2)$, for all $r \in \Sigma$.

The following characterization is straightforward to prove and can be found in [7].

Lemma 22. For all pointed interpretations (\mathcal{I}_1, d_1) and (\mathcal{I}_2, d_2) , all finite signatures Σ , and all $m \geq 0$: $(\mathcal{I}_1, d_1) \equiv_{\Sigma}^m (\mathcal{I}_2, d_2)$ if, and only if, $(\mathcal{I}_1, d_1) \sim_{\Sigma}^m (\mathcal{I}_2, d_2)$.

Lemma 23. Let (\mathcal{I}_1, d_1) and (\mathcal{I}_2, d_2) be pointed Σ -interpretations such that $(\mathcal{I}_1, d_1) \sim_{\Sigma}^m (\mathcal{I}_2, d_2)$. Then there exist Σ -tree interpretations \mathcal{J}_1 and \mathcal{J}_2 such that

- $(\mathcal{J}_1, \rho^{\mathcal{J}_1}) \sim_{\Sigma} (\mathcal{I}_1, d_1)$;
- $(\mathcal{J}_2, \rho^{\mathcal{J}_2}) \sim_{\Sigma} (\mathcal{I}_2, d_2)$;

- $\mathcal{J}_1^{\leq m} = \mathcal{J}_2^{\leq m}$.

Moreover, if \mathcal{I}_1 and \mathcal{I}_2 have finite outdegree, then one can find such \mathcal{J}_1 and \mathcal{J}_2 that have finite outdegree.

Proof. Assume (\mathcal{I}_1, d_1) and (\mathcal{I}_2, d_2) are given. For $j = 1, 2$, we define \mathcal{J}_j as follows. The domain $\Delta^{\mathcal{J}_j}$ consists of all words

$$v_0 r_0 v_1 r_1 \cdots v_n$$

such that

- $v_0 = (d_1, d_2)$;
- for all $i \leq m$ there are e_i, f_i such that $v_i = (e_i, f_i)$ and $(\mathcal{I}_1, e_i) \sim_{\Sigma}^{m-i} (\mathcal{I}_2, f_i)$;
- for all $i < m$: if $v_i = (e_i, f_i)$ and $v_{i+1} = (e_{i+1}, f_{i+1})$, then $(e_i, e_{i+1}) \in r_i^{\mathcal{I}_1}$ and $(f_i, f_{i+1}) \in r_i^{\mathcal{I}_2}$;
- for all $i > m$: $v_i \in \Delta^{\mathcal{J}_j}$;
- for all $v_m = (e_m, f_m)$: if $j = 1$ then $(e_m, v_{m+1}) \in r_m^{\mathcal{I}_1}$; and if $j = 2$ then $(f_m, v_{m+1}) \in r_m^{\mathcal{I}_2}$;
- for all $i > m$: $(v_i, v_{i+1}) \in r_i^{\mathcal{J}_j}$.

For all concept names A , we set

$$A^{\mathcal{J}_j} = \{v_0 \cdots v_n \in \Delta^{\mathcal{J}_j} \mid n > m, v_n \in A^{\mathcal{I}_j}\} \cup \{(e_0, f_0) \cdots (e_n, f_n) \in \Delta^{\mathcal{J}_j} \mid n \leq m, e_n \in A^{\mathcal{I}_1}\},$$

and for all role names r , we define

$$r^{\mathcal{J}_j} = \{(w, wrv) \mid w, wrv \in \Delta^{\mathcal{J}_j}, v \in (\Delta^{\mathcal{I}_j} \cup \Delta^{\mathcal{I}_1} \times \Delta^{\mathcal{I}_2})\}$$

It is straightforward to prove that

$$S_1 = \{(e, w(e, f)) \mid e \in \Delta^{\mathcal{I}_1}, w(e, f) \in \Delta^{\mathcal{J}_1}\} \cup \{(e, we) \mid e \in \Delta^{\mathcal{I}_1}, we \in \Delta^{\mathcal{J}_1}\}$$

is a Σ -bisimulation between $(\mathcal{I}_1, \rho^{\mathcal{I}_1})$ and $(\mathcal{J}_1, \rho^{\mathcal{J}_1})$. A Σ -bisimulation S_2 between $(\mathcal{I}_2, \rho^{\mathcal{I}_2})$ and $(\mathcal{J}_2, \rho^{\mathcal{J}_2})$ can be constructed in the same way. Clearly, $\mathcal{J}_1^{\leq m} = \mathcal{J}_2^{\leq m}$. \square

Theorem 9. Let \mathcal{T} be a TBox, $\Sigma \subseteq \text{sig}(\mathcal{T})$ a signature, and $m \geq 0$. Then $\mathcal{T}_{\Sigma, m}$ is not a uniform Σ -interpolant of \mathcal{T} iff

($*_m$) there exist two Σ -tree interpretations, \mathcal{I}_1 and \mathcal{I}_2 , of finite outdegree such that

1. $\mathcal{I}_1^{\leq m} = \mathcal{I}_2^{\leq m}$;
2. $(\mathcal{I}_1, \rho^{\mathcal{I}_1}) \models \exists_{\Sigma}^{\sim} \mathcal{T}$;
3. $(\mathcal{I}_2, \rho^{\mathcal{I}_2}) \not\models \exists_{\Sigma}^{\sim} \mathcal{T}$;
4. For all sons d of $\rho^{\mathcal{I}_2}$: $(\mathcal{I}_2, d) \models \exists_{\Sigma}^{\sim} \mathcal{T}$.

Proof. Assume first that $\mathcal{T}_{\Sigma, m}$ is not a uniform Σ -interpolant of \mathcal{T} . We show ($*_m$). There exists $m' \geq m$ such that $\mathcal{T}_{\Sigma, m'} \not\models \mathcal{T}_{\Sigma, m'+1}$. There exists a Σ -tree interpretation \mathcal{I} of finite outdegree such that $\mathcal{I} \models \mathcal{T}_{\Sigma, m'}$ and $\rho^{\mathcal{I}} \notin C^{\mathcal{I}}$ for some C with $\top \sqsubseteq C \in \mathcal{T}_{\Sigma, m'+1}$.

As $\mathcal{I} \models \mathcal{T}_{\Sigma, m'}$, the concept

$$D = \bigsqcap_{E \in C_f^{m'}(\Sigma), \rho^{\mathcal{I}} \in E^{\mathcal{I}}} E$$

is satisfiable w.r.t. \mathcal{T} . There exists a Σ -tree interpretation \mathcal{J} of finite outdegree that is a model of \mathcal{T} such that $(\mathcal{I}, \rho^{\mathcal{I}}) \equiv_{\Sigma}^{m'} (\mathcal{J}, \rho^{\mathcal{J}})$. By Lemma 22, we have $(\mathcal{I}, \rho^{\mathcal{I}}) \sim_{\Sigma}^{m'} (\mathcal{J}, \rho^{\mathcal{J}})$. By Lemma 23, and closure under composition of (m)-bisimulations, we can assume that \mathcal{J} is a Σ -tree interpretation of finite outdegree with

- $\mathcal{I}^{\leq m'} = \mathcal{J}^{\leq m'}$.
- $(\mathcal{J}, \rho^{\mathcal{J}}) \models \exists_{\Sigma}^{\sim} \mathcal{T}$.
- $\mathcal{I} \models \mathcal{T}_{\Sigma, m'}$.
- $\rho^{\mathcal{I}} \notin C^{\mathcal{I}}$ for some $(\top \sqsubseteq C) \in \mathcal{T}_{\Sigma, m'+1}$.

For every son d of $\rho^{\mathcal{I}}$, as $\mathcal{I}(d)$ is a model of $\mathcal{T}_{\Sigma, m'}$ we can argue as above and find a Σ -tree interpretation \mathcal{K}_d of finite outdegree such that

$$(\mathcal{I}(d), d) \sim_{\Sigma}^{m'} (\mathcal{K}_d, \rho^{\mathcal{K}_d}) \models \exists_{\Sigma}^{\sim} \mathcal{T}.$$

We have

$$(\mathcal{J}(d), d) \sim_{\Sigma}^{m'-1} (\mathcal{K}_d, \rho^{\mathcal{K}_d}).$$

Thus, by Lemma 23, we find Σ -tree interpretations \mathcal{J}_d and \mathcal{M}_d of finite outdegree such that

- $(\mathcal{J}_d, \rho^{\mathcal{J}_d}) \sim_{\Sigma} (\mathcal{J}(d), d)$;
- $(\mathcal{K}_d, \rho^{\mathcal{K}_d}) \sim_{\Sigma} (\mathcal{M}_d, \rho^{\mathcal{M}_d})$;
- $\mathcal{J}_d^{\leq m'-1} = \mathcal{M}_d^{\leq m'-1}$.

Now define \mathcal{I}_1 by replacing, for every son d of $\rho^{\mathcal{J}}$, $\mathcal{J}(d)$ by \mathcal{J}_d in \mathcal{J} . Define \mathcal{I}_2 by replacing, for every son d of $\rho^{\mathcal{I}}$, $\mathcal{I}(d)$ by \mathcal{M}_d in \mathcal{I} . It is readily checked that \mathcal{I}_1 and \mathcal{I}_2 are as required:

- $\mathcal{I}_1^{\leq m} = \mathcal{I}_2^{\leq m}$: since $m' \geq m$ it is sufficient to show $\mathcal{I}_1^{\leq m'} = \mathcal{I}_2^{\leq m'}$. But since $\mathcal{I}^{\leq 1} = \mathcal{J}^{\leq 1}$ this follows from $\mathcal{J}_d^{\leq m'-1} = \mathcal{M}_d^{\leq m'-1}$ for every son d of $\rho^{\mathcal{I}}$.
- $(\mathcal{I}_1, \rho^{\mathcal{I}_1}) \models \exists_{\Sigma}^{\sim} \mathcal{T}$ follows from $(\mathcal{J}, \rho^{\mathcal{J}}) \models \exists_{\Sigma}^{\sim} \mathcal{T}$ and $(\mathcal{I}_1, \rho^{\mathcal{I}_1}) \sim_{\Sigma} (\mathcal{J}, \rho^{\mathcal{J}})$.
- $(\mathcal{I}_2, \rho^{\mathcal{I}_2}) \not\models \exists_{\Sigma}^{\sim} \mathcal{T}$ follows if $\rho^{\mathcal{I}_2} \notin C^{\mathcal{I}_2}$ for some Σ -concept C such that $\mathcal{T} \models \top \sqsubseteq C$. By construction, there exists $C \in C_f^{m'+1}(\Sigma)$ such that $\rho^{\mathcal{I}} \notin C^{\mathcal{I}}$ and $(\top \sqsubseteq C) \in \mathcal{T}_{\Sigma, m'+1}$. Thus, by Lemma 22, it is sufficient to show $(\mathcal{I}_2, \rho^{\mathcal{I}_2}) \sim_{\Sigma}^{m'+1} (\mathcal{I}, \rho^{\mathcal{I}})$. This follows if $(\mathcal{I}(d), d) \sim_{\Sigma}^{m'} (\mathcal{M}_d, \rho^{\mathcal{M}_d})$, for every son d of $\rho^{\mathcal{I}}$. But this follows from

$$(\mathcal{I}(d), d) \sim_{\Sigma}^{m'} (\mathcal{K}_d, \rho^{\mathcal{K}_d}) \sim_{\Sigma} (\mathcal{M}_d, \rho^{\mathcal{M}_d}).$$

- For all sons d of $\rho^{\mathcal{I}_2}$: $(\mathcal{I}_2, d) \models \exists_{\Sigma}^{\sim} \mathcal{T}$. This follows from

$$(\mathcal{I}_2(d), d) = (\mathcal{M}_d, \rho^{\mathcal{M}_d}) \sim_{\Sigma} (\mathcal{K}_d, \rho^{\mathcal{K}_d}) \models \exists_{\Sigma}^{\sim} \mathcal{T}.$$

Now assume that \mathcal{I}_1 and \mathcal{I}_2 satisfy ($*_m$). Then \mathcal{I}_2 is a model of $\mathcal{T}_{\Sigma, m}$. For assume this is not the case. Then $d \notin C^{\mathcal{I}_2}$ for some C with $\top \sqsubseteq C \in \mathcal{T}_{\Sigma, m}$. If $d = \rho^{\mathcal{I}_2}$, then $d \notin C^{\mathcal{I}_1}$ because of Point 1. This contradicts $\mathcal{I}_1 \models \exists_{\Sigma}^{\sim} \mathcal{T}$. If $d \neq \rho^{\mathcal{I}_2}$, then $d \in \mathcal{I}(d')$ for some son d' of $\rho^{\mathcal{I}_2}$. Hence $d \notin C^{\mathcal{I}(d')}$, which contradicts Point 4.

Now assume that $\mathcal{T}_{\Sigma, m}$ is a uniform Σ -interpolant of \mathcal{T} . As \mathcal{I}_2 is a model of $\mathcal{T}_{\Sigma, m}$ and has finite outdegree, we obtain from the modified version of Theorem 6 proved above that $(\mathcal{I}_2, \rho^{\mathcal{I}_2}) \models \exists_{\Sigma}^{\approx} \mathcal{T}$, which contradicts Point 3. \square

We show Example 3 (iv): for \mathcal{T}_4 consisting of

1. $A \sqsubseteq \exists r.B$;
2. $A_0 \sqsubseteq \exists r.(A_1 \sqcap B)$;
3. $E \equiv A_1 \sqcap B \sqcap \exists r.(A_2 \sqcap B)$;

and $\Sigma_4 = \{A, r, A_0, A_1, E\}$, there is no uniform Σ_4 -interpolant of \mathcal{T}_4 .

Proof. It is sufficient to show $(*_m)$ for all $m > 0$. Let $\mathcal{I}_1 = (\{0, \dots, m+1, (a, 2), \dots, (a, m+1)\}, \cdot^{\mathcal{I}_1})$, where

$$\begin{aligned} A^{\mathcal{I}_1} &= \{1, \dots, m\} \\ E^{\mathcal{I}_1} &= \emptyset \\ r^{\mathcal{I}_1} &= \{(i, i+1) \mid 0 \leq i \leq m\} \cup \\ &\quad \{(i, (a, i+1)) \mid 1 \leq i \leq m\} \\ A_1^{\mathcal{I}_1} &= \{1, \dots, m+1\} \\ A_2^{\mathcal{I}_1} &= \{(a, 2), \dots, (a, m+1)\} \\ A_0^{\mathcal{I}_1} &= \{0\} \end{aligned}$$

Then $(\mathcal{I}_1, 0) \models \exists_{\Sigma_4}^{\approx} \mathcal{T}_4$ because the expansion of \mathcal{I}_1 by $B^{\mathcal{I}_1} = \{1, \dots, m+1\}$ is a model of \mathcal{T}_4 .

Define \mathcal{I}_2 as the restriction of \mathcal{I}_1 to $\Delta^{\mathcal{I}_1} \setminus \{m+1\}$. By definition, $\mathcal{I}_1^{\leq m} = \mathcal{I}_2^{\leq m}$,

Claim 1. $(\mathcal{I}_2, 0) \not\models \exists_{\Sigma_4}^{\approx} \mathcal{T}_4$.

Assume $(\mathcal{I}_2, 0) \models \exists_{\Sigma_4}^{\approx} \mathcal{T}_4$. Take $(\mathcal{I}_2, 0) \sim_{\Sigma} (\mathcal{J}, 0')$ with \mathcal{J} a model of \mathcal{T}_4 . Let S be the Σ -bisimulation with $(0, 0') \in S$. By inclusion (2.) there exists $1'$ with $(1, 1') \in S$ such that $1' \in (\neg E \sqcap A \sqcap A_1 \sqcap B)^{\mathcal{J}}$. By inclusion (1.) there exists an r -successor $2'$ of $1'$ that is in $B^{\mathcal{J}}$. We have $(2, 2') \in S$ or $((a, 2), 2') \in S$. But $((a, 2), 2') \notin S$ because otherwise $2' \in (A_2 \sqcap B)^{\mathcal{J}}$ which, since $1' \in (A_1 \sqcap B)^{\mathcal{J}}$, would imply, by inclusion (3.), that $1' \in E^{\mathcal{J}}$, a contradiction. Thus, $(2, 2') \in S$ and so $2' \in (\neg E \sqcap A \sqcap A_1 \sqcap B)^{\mathcal{J}}$. One can now show in same way by induction that there is a m' with $(m, m') \in S$ such that $m' \in (\neg E \sqcap A \sqcap A_1 \sqcap B)^{\mathcal{J}}$. All r -successors of m' are in $A_2^{\mathcal{J}}$ since all r -successors of m are in $A_2^{\mathcal{J}}$. By inclusions (2.) and (3.) and since $m' \notin E^{\mathcal{J}}$ this leads to a contradiction.

As 1 is the only r -successor of 0 in \mathcal{I}_2 , it remains to show that $(\mathcal{I}_2, 1) \models \exists_{\Sigma_4}^{\approx} \mathcal{T}_4$. Let \mathcal{I}'_2 be the restriction of \mathcal{I}_2 to $\Delta^{\mathcal{I}_2} \setminus \{0\}$. Then $(\mathcal{I}_2, 1) \models \exists_{\Sigma_4}^{\approx} \mathcal{T}_4$ follows from the observation that the expansion of \mathcal{I}'_2 by setting $B^{\mathcal{I}'_2} = \{(a, 2), \dots, (a, m+1)\}$ is a model of \mathcal{T}_4 . \square

In a tree interpretation \mathcal{I} , we set $\text{dist}(\rho^{\mathcal{I}}, d) = k$ and say that the *depth* of d in \mathcal{I} is k iff d can be reached from $\rho^{\mathcal{I}}$ in exactly k steps.

Example 11 Let \mathcal{T} be a TBox and Σ a signature such that $\text{sig}(\mathcal{T}) \setminus \Sigma$ consists of stratified concept names only, i.e., we

want to *forget* a set of stratified concept names. Then the existence of a uniform Σ -interpolant of \mathcal{T} is guaranteed; moreover, $\mathcal{T}_{\Sigma, m}$ is such an interpolant, where $m = \max\{\text{rd}(C) \mid C \in \text{conc}(\mathcal{T})\}$.

Proof. Assume $\mathcal{I}_1, \mathcal{I}_2$ satisfy $(*_m)$. There exists a tree-interpretation \mathcal{J}_1 that is a model of \mathcal{T} such that $(\mathcal{I}_1, \rho^{\mathcal{I}_1}) \sim_{\Sigma} (\mathcal{J}_1, \rho^{\mathcal{J}_1})$, and for every r -successor d_r of $\rho^{\mathcal{I}_2}$ there exists a tree-interpretation \mathcal{J}_{d_r} that is a model of \mathcal{T} with $(\mathcal{I}_2, d_r) \sim_{\Sigma} (\mathcal{J}_{d_r}, \rho_{d_r})$. We may assume that \mathcal{I}_1 is the Σ -reduct of \mathcal{J}_1 and that every $\mathcal{I}_2(d_r)$ coincides with the Σ -reduct of \mathcal{J}_{d_r} . Now expand \mathcal{I}_2 to an interpretation \mathcal{J}_2 as follows: for every $B \in \text{sig}(\mathcal{T}) \setminus \Sigma$ of level $k \leq m$ set

$$\begin{aligned} B^{\mathcal{J}_2} &= \{d \mid d \in B^{\mathcal{J}_1} \wedge \text{dist}(\rho^{\mathcal{I}_2}, d) = k\} \cup \\ &\quad \{d \mid d \in B^{\mathcal{J}_{d_r}} \wedge \text{dist}(\rho^{\mathcal{I}_2}, d) \neq k\}. \end{aligned}$$

We show that \mathcal{J}_2 is a model of \mathcal{T} ; and have derived a contradiction as $\mathcal{I}_2 \not\models \exists_{\Sigma}^{\approx} \mathcal{T}$. Let $C \sqsubseteq D \in \mathcal{T}$ and assume that $d \in C^{\mathcal{J}_2} \setminus D^{\mathcal{J}_2}$. Let $\text{dist}(\rho^{\mathcal{J}_2}, d) = l$. If $l = 0$, then $d \in X^{\mathcal{J}_2}$ iff $d \in X^{\mathcal{I}_1}$ for all concepts X in \mathcal{T} , by the definition of the expansion. Thus, $d \in C^{\mathcal{I}_1} \setminus D^{\mathcal{I}_1}$ which contradicts that \mathcal{I}_1 is a model of \mathcal{T} . If $l > 0$, then d is in the domain of some $\mathcal{I}_2(d_r)$. Then $d \in X^{\mathcal{J}_2}$ iff $d \in X^{\mathcal{J}_{d_r}}$ for all concepts X in \mathcal{T} , by the definition of the expansion. Thus, $d \in C^{\mathcal{J}_{d_r}} \setminus D^{\mathcal{J}_{d_r}}$ which contradicts that \mathcal{J}_{d_r} is a model of \mathcal{T} . \square

Fix a TBox \mathcal{T} and $\Sigma \subseteq \text{sig}(\mathcal{T})$. Set $(\mathcal{I}_1, d_1) \sim_e (\mathcal{I}_2, d_2)$ iff $\text{Ext}^{\mathcal{I}_1}(d_1) = \text{Ext}^{\mathcal{I}_2}(d_2)$. Note that the number of \sim_e -equivalence classes is bounded by $M_{\mathcal{T}}$.

Lemma 24. Let \mathcal{I} be a tree interpretation and $d \in \Delta^{\mathcal{I}}$. Assume $(\mathcal{I}(d), d) \sim^e (\mathcal{J}, \rho^{\mathcal{J}})$ for a tree interpretation \mathcal{J} . Replace $\mathcal{I}(d)$ by \mathcal{J} in \mathcal{I} and denote the resulting tree interpretation by \mathcal{K} . Then $\mathcal{I} \models \exists_{\Sigma}^{\approx} \mathcal{T}$ if, and only if, $\mathcal{K} \models \exists_{\Sigma}^{\approx} \mathcal{T}$.

Proof. Let \mathcal{I} , d , \mathcal{J} , and \mathcal{K} be as in the formulation of Lemma 24. Assume $\mathcal{I} \models \exists_{\Sigma}^{\approx} \mathcal{T}$. There exists a tree-interpretation \mathcal{I}' that is a model of \mathcal{T} such that $(\mathcal{I}, \rho^{\mathcal{I}}) \sim_{\Sigma} (\mathcal{I}', \rho^{\mathcal{I}'})$. We may assume that there is a Σ -bisimulation S between $(\mathcal{I}, \rho^{\mathcal{I}})$ and $(\mathcal{I}', \rho^{\mathcal{I}'})$ such that S^- is an injective relation and such that $(e, e') \in S$ implies that e is reached from $\rho^{\mathcal{I}}$ along the same path as e' from $\rho^{\mathcal{I}'}$. Let $S(d) = \{d' \mid (d, d') \in S\}$. Consider, for every $d' \in S(d)$, a tree-interpretation $\mathcal{K}_{d'}$ satisfying \mathcal{T} such that

- $\text{tp}^{\mathcal{I}', \mathcal{T}}(d') = \text{tp}^{\mathcal{K}_{d'}, \mathcal{T}}(\rho^{\mathcal{K}_{d'}})$;
- $(\mathcal{J}, \rho^{\mathcal{J}}) \sim_{\Sigma} (\mathcal{K}_{d'}, \rho^{\mathcal{K}_{d'}})$.

Such interpretations $\mathcal{K}_{d'}$ exist by the definition of the equivalence relation \sim^e . Now replace, in \mathcal{I}' and for all $d' \in S(d)$, the tree interpretation $\mathcal{I}'(d')$ by $\mathcal{K}_{d'}$, and denote the resulting interpretation by \mathcal{K}' . \mathcal{K}' is a model of \mathcal{T} since $\text{tp}^{\mathcal{I}', \mathcal{T}}(d') = \text{tp}^{\mathcal{K}_{d'}, \mathcal{T}}(\rho^{\mathcal{K}_{d'}})$, \mathcal{I}' is a model of \mathcal{T} and all $\mathcal{K}_{d'}$ are models of \mathcal{T} . It remains to show that $(\mathcal{K}, \rho^{\mathcal{K}}) \sim_{\Sigma} (\mathcal{K}', \rho^{\mathcal{K}'})$. Take for every $d' \in S(d)$ a Σ -bisimulation $S_{d'}$ between $(\mathcal{J}, \rho^{\mathcal{J}})$ and $(\mathcal{K}_{d'}, \rho^{\mathcal{K}_{d'}})$. Let S' be the restriction of S to

$$(\Delta^{\mathcal{I}} \setminus \Delta^{\mathcal{I}(d)}) \times (\Delta^{\mathcal{I}'} \setminus (\bigcup_{d' \in S(d)} \Delta^{\mathcal{I}'(d')}))$$

It is not difficult to show that $S' \cup \bigcup_{d' \in S(d)} S_{d'}$ is the required Σ -bisimulation between $(\mathcal{K}, \rho^{\mathcal{K}})$ and $(\mathcal{K}', \rho^{\mathcal{K}'})$. \square

Theorem 12. Let \mathcal{T} be a TBox and $\Sigma \subseteq \text{sig}(\mathcal{T})$. Then there does not exist a uniform Σ -interpolant of \mathcal{T} iff $(*_M^{\mathcal{T}+1})$ from Theorem 9 holds, where $M_{\mathcal{T}} := 2^{2^{|\mathcal{T}|}}$.

Proof. By Theorem 9, it is sufficient to prove that $(*_M^{\mathcal{T}+1})$ implies $(*_m)$ for all $m \geq M_{\mathcal{T}}^2 + 1$. Take Σ -tree interpretations \mathcal{I}_1 and \mathcal{I}_2 satisfying $(*_m)$ of Theorem 9 for some $m \geq M_{\mathcal{T}}^2 + 1$. We show that there exist Σ -tree interpretations \mathcal{J}_1 and \mathcal{J}_2 satisfying $(*_m)$ of Theorem 9. The implication then follows by induction.

Let D be the set of $d \in \Delta^{\mathcal{I}_1}$ with $\text{dist}(\rho^{\mathcal{I}_1}, d) = m$ such that $\mathcal{I}_1(d)^{\leq 1} \neq \mathcal{I}_2(d)^{\leq 1}$ (i.e., the restrictions of \mathcal{I}_1 to $\{d' \mid d' \text{ son of } d \text{ in } \mathcal{I}_1\}$ and \mathcal{I}_2 to $\{d' \mid d' \text{ son of } d \text{ in } \mathcal{I}_2\}$ do not coincide). If $D = \emptyset$, then $\mathcal{I}_1^{\leq m+1} = \mathcal{I}_2^{\leq m+1}$ and the claim is proved. Otherwise choose $f \in D$ and consider the path

$$\rho^{\mathcal{I}_1} = d_0 r_0 d_1 \cdots r_{m-1} d_m = f$$

with $(d_i, d_{i+1}) \in r_i^{\mathcal{I}_1}$ for all $i < m$. As $m \geq M_{\mathcal{T}}^2 + 1$, there exists $0 < i < j \leq m$ such that both,

$$(\mathcal{I}_1, d_i) \sim^e (\mathcal{I}_1, d_j), \quad (\mathcal{I}_2, d_i) \sim^e (\mathcal{I}_2, d_j).$$

Replace $\mathcal{I}_1(d_j)$ by $\mathcal{I}_1(d_i)$ in \mathcal{I}_1 and denote the resulting interpretation by \mathcal{K}_1 . Similarly, replace $\mathcal{I}_2(d_j)$ by $\mathcal{I}_2(d_i)$ in \mathcal{I}_2 and denote the resulting interpretation \mathcal{K}_2 . By Lemma 24, \mathcal{K}_1 and \mathcal{K}_2 still have Properties (1)–(4). Moreover, the set D' of all $d \in \Delta^{\mathcal{I}_1}$ with $\text{dist}(\rho^{\mathcal{I}_1}, d) = m$ such that $\mathcal{K}_1(d)^{\leq 1} \neq \mathcal{K}_2(d)^{\leq 1}$ is a subset of D not containing f . Thus, we can proceed with D' in the same way as above until the set is empty. Denote the resulting interpretations by \mathcal{J}_1 and \mathcal{J}_2 , respectively. They still have Properties (1)–(4), but now for some $m' > m$. \square

9 Proofs for Section 4

We start with establishing some basic results about APTAs. To formulate and prove these, we make some technicalities more formal than in the main paper. Recall that a run is a pair (T, ℓ) with T a tree. Let us make precise what exactly we mean by ‘tree’ here. A *tree* is a non-empty (finite or infinite) prefix-closed subset $T \subseteq S^*$, for some set S . If $d \in T$ and $d \cdot c \in T$ with $d \in S^*$ and $c \in S$, then the node $d \cdot c$ is a *son* of the node d in T . A node $d \in T$ that has no sons is a *leaf*. We measure the size of an APTA primarily in the number of states. To define a more fine-grained measure, we use $\|\mathcal{A}\|$ to denote the *size* of \mathcal{A} , i.e., $|Q| + |\Sigma_N| + |\Sigma_E|$. Note that the size of (the representation of) all other components of the automaton is bounded polynomially in $\|\mathcal{A}\|$. In particular, we can w.l.o.g. assume that the values in Ω are bounded by $2\|Q\|$.

Lemma 25. Let $\mathcal{A}_i = (Q_i, \Sigma_N, \Sigma_E, q_{0,i}, \delta_i, \Omega_i)$ be APTAs, $i \in \{1, 2\}$ with $Q_1 \cap Q_2 = \emptyset$. Then there is an APTA

1. $\mathcal{A}' = (Q_1, \Sigma_N, \Sigma_E, q_{0,1}, \delta', \Omega')$ such that $L(\mathcal{A}') = \overline{L(\mathcal{A}_1)}$;
2. $\mathcal{A}'' = (Q_1 \uplus Q_2 \uplus \{q_0\}, \Sigma_N, \Sigma_E, q_0, \delta'', \Omega'')$ such that $L(\mathcal{A}'') = L(\mathcal{A}_1) \cap L(\mathcal{A}_2)$.

Moreover, \mathcal{A}' and \mathcal{A}'' can be constructed in time $p(\|\mathcal{A}\|)$, p a polynomial.

Proof. (sketch) The construction of \mathcal{A}' is based on the standard dualization construction first given in [3], i.e., δ' is obtained from δ by swapping true and false, A and $\neg A$, \wedge and \vee , and diamonds and boxes, and setting $\Omega'(q) = \Omega(q) + 1$ for all $q \in Q$. The construction of \mathcal{A}'' is standard as well: add a fresh initial state q_0 with $\delta(q_0) = q_{0,1} \wedge q_{0,2}$, and define δ'' and Ω'' as the fusion of the respective components of \mathcal{A}_1 and \mathcal{A}_2 (e.g., $\delta''(q) = \delta_1(q)$ for all $q \in Q_1$ and $\delta''(q) = \delta_2(q)$ for all $q \in Q_2$). \square

The following lemma shows that a language accepted by an APTA is closed under bisimulation. It implies that whenever for an APTA \mathcal{A} we have $L(\mathcal{A}) \neq \emptyset$, then there is a pointed tree interpretation (\mathcal{I}, d) with $(\mathcal{I}, d) \in L(\mathcal{A})$. As a notational convention, whenever x is a node in a $Q \times \Delta^{\mathcal{I}}$ -labelled tree and $\ell(x) = (q, d)$, then we use $\ell_1(x)$ to denote q and $\ell_2(x)$ to denote d .

Lemma 26. Let $\mathcal{A} = (Q, \Sigma_N, \Sigma_E, q_0, \delta, \Omega)$ be an APTA, $(\mathcal{I}, d) \in L(\mathcal{A})$, and $(\mathcal{I}, d) \sim_{\Sigma_N \cup \Sigma_E} (\mathcal{J}, e)$. Then $(\mathcal{J}, e) \in L(\mathcal{A})$.

Proof. Let $(\mathcal{I}, d) \in L(\mathcal{A})$, and $(\mathcal{I}, d) \sim_{\Sigma_N \cup \Sigma_E} (\mathcal{J}, e)$. Moreover, let (T, ℓ) be an accepting run of \mathcal{A} on (\mathcal{I}, d) . We inductively construct a $Q \times \Delta^{\mathcal{J}}$ -labelled tree (T', ℓ') , along with a map $\mu : T' \rightarrow T$ such that $\mu(y) = x$ implies $\ell_1(x) = \ell'_1(y)$ and $(\mathcal{I}, \ell_2(x)) \sim_{\Sigma_N \cup \Sigma_E} (\mathcal{J}, \ell'_2(y))$:

- start with $T' = \{\varepsilon\}$, $\ell'(\varepsilon) = (q, e)$, and $\mu(\varepsilon) = \varepsilon$;
- if $y \in T'$ is a leaf, $\ell'_1(y) = q' \wedge q''$, and $\mu(y) = x$, then there are sons x', x'' of x with $\ell(x') = (q', \ell_2(x))$ and $\ell(x'') = (q'', \ell_2(x))$; add fresh $y \cdot c'$ and $y \cdot c''$ to T' and put $\ell'(y \cdot c') = (q', \ell'_2(y))$, $\ell'(y \cdot c'') = (q'', \ell'_2(y))$, $\mu(y \cdot c') = x'$, and $\mu(y \cdot c'') = x''$;
- if $y \in T'$ is a leaf, $\ell'_1(y) = q' \vee q''$, and $\mu(y) = x$, then there is a son x' of x with $\ell_1(x') \in \{q', q''\}$ and $\ell_2(x') = \ell_2(x)$; add a fresh $y \cdot c'$ to T' and put $\ell'(y \cdot c') = (\ell_1(x'), \ell'_2(y))$ and $\mu(y \cdot c') = x'$;
- if $y \in T'$ is a leaf, $\ell'_1(y) = \langle r \rangle q'$, and $\mu(y) = x$, then there is an $(\ell_2(x), d) \in r^{\mathcal{I}}$ and a son x' of x with $\ell(x') = (q', d)$; since $(\mathcal{I}, \ell_2(x)) \sim_{\Sigma_N \cup \Sigma_E} (\mathcal{J}, \ell'_2(y))$, there is an $(\ell'_2(y), d')$ with $(\mathcal{I}, d) \sim_{\Sigma_N \cup \Sigma_E} (\mathcal{J}, d')$; add a fresh $y \cdot c'$ to T' and put $\ell'(y \cdot c') = (q', d')$ and $\mu(y \cdot c') = x'$;
- if $y \in T'$ is a leaf, $\ell'_1(y) = [r]q'$, and $\mu(y) = x$, then do the following for every $(\ell'_2(y), d') \in r^{\mathcal{J}}$: since $(\mathcal{I}, \ell_2(x)) \sim_{\Sigma_N \cup \Sigma_E} (\mathcal{J}, \ell'_2(y))$, there is an $(\ell_2(x), d) \in r^{\mathcal{I}}$ with $(\mathcal{I}, d) \sim_{\Sigma_N \cup \Sigma_E} (\mathcal{J}, d')$, and thus also a son x' of x with $\ell(x') = (q', d)$; add a fresh $y \cdot c'$ to T' and put $\ell'(y \cdot c') = (q', d')$ and $\mu(y \cdot c') = x'$.

It can be verified that (T', ℓ') is an accepting run of \mathcal{A} on (\mathcal{J}, e) . \square

Finally, we fix the complexity of the emptiness problem of APTAs. The following is proved in [20] using a reduction to parity games.

Theorem 27 (Wilke). Let $\mathcal{A} = (Q, \Sigma_N, \Sigma_E, q_0, \delta, \Omega)$ be an APTA. Then the emptiness of $L(\mathcal{A})$ can be decided in time $2^{p(\|\mathcal{A}\|)}$, p a polynomial.

The following lemma establishes the correctness of the construction of the automata $\mathcal{A}_{\mathcal{T}_\Sigma}$ for Theorem 15.

Lemma 28. $(\mathcal{I}, d_0) \in L(\mathcal{A}_{\mathcal{T}, \Sigma})$ iff $(\mathcal{I}, d_0) \models \exists_{\Sigma}^{\approx} \mathcal{T}$, for all pointed Σ -interpretations (\mathcal{I}, d_0) .

Proof. Let (\mathcal{I}, d_0) be a pointed Σ -interpretation. By definition of $\mathcal{A}_{\mathcal{T}, \Sigma}$, we have $(\mathcal{I}, d_0) \in L(\mathcal{A}_{\mathcal{T}, \Sigma})$ iff there exists a $Q \times \Delta^{\mathcal{I}}$ -labelled tree (T, ℓ) such that

1. $\ell(\varepsilon) = (q_0, d_0)$;
2. there exists a son x of ε and $t \in \text{TP}(\mathcal{T})$ such that $\ell(x) = (t, d_0)$;
3. if $\ell(x) = (t, d)$, then
 - (a) $d \in A^{\mathcal{I}}$ for all $A \in t \cap \mathbf{N}_C \cap \Sigma$;
 - (b) $d \notin A^{\mathcal{I}}$ for all $A \in (\mathbf{N}_C \cap \Sigma) \setminus t$;
 - (c) for all $r \in \Sigma$ and $(d, d') \in r^{\mathcal{I}}$, there exist $t' \in \text{TP}(\mathcal{T})$ such that $t \rightsquigarrow_r t'$ and a son y of x such that $\ell(y) = (t', d')$;
 - (d) for all $\exists r.C \in t$ with $r \in \Sigma$, there exists $(d, d') \in r^{\mathcal{I}}$ and $t' \in \text{TP}(\mathcal{T})$ such that $C \in t'$ and $t \rightsquigarrow_r t'$, and $\ell(y) = (t', d')$ for some son y of x .

Assume that $(\mathcal{I}, d_0) \models \exists_{\Sigma}^{\approx} \mathcal{T}$ and let (\mathcal{J}, e_0) be a pointed model of \mathcal{T} such that $(\mathcal{I}, d_0) \sim_{\Sigma} (\mathcal{J}, e_0)$. A path is a sequence $(d_1, e_1) \cdots (d_n, e_n)$, $n \geq 0$, with $d_1, \dots, d_n \in \Delta^{\mathcal{I}}$ and $e_1, \dots, e_n \in \Delta^{\mathcal{J}}$ such that

- (i) $d_1 = d_0$;
- (ii) $(d_i, d_{i+1}) \in r^{\mathcal{I}}$ for some $r \in \Sigma \cap \mathbf{N}_R$, for $1 \leq i < n$;
- (iii) $(\mathcal{I}, d_i) \sim_{\Sigma} (\mathcal{J}, e_i)$ for $1 \leq i \leq n$.

Define a $Q \times \Delta^{\mathcal{I}}$ -labelled tree (T, ℓ) by setting

- T to the set of all paths;
- $\ell(\varepsilon) = (q_0, d_0)$;
- $\ell((d_1, e_1) \cdots (d_n, e_n)) = (\text{tp}^{\mathcal{J}}(e_n), d_n)$ for all paths $(d_1, e_1) \cdots (d_n, e_n) \neq \varepsilon$.

One can now verify that (T, ℓ) satisfies Conditions 1 to 3, thus $(\mathcal{I}, d_0) \in L(\mathcal{A}_{\mathcal{T}, \Sigma})$. In fact, Conditions 1 and 2 are immediate and Conditions 3a and 3b are a consequence of the definition of ℓ and (iii). As for Condition 3c, let $\ell(x) = (t, d)$, $(d, d') \in r^{\mathcal{I}}$, and $r \in \Sigma$, and assume $x = (d_1, e_1) \cdots (d_n, e_n)$. Then $d = d_n$, and $d_n \sim_{\Sigma} e_n$ and $(d, d') \in r^{\mathcal{I}}$ yield an (e_n, e') with $d' \sim_{\Sigma} e'$. Thus $y := (d_1, e_1) \cdots (d_n, e_n)(d', e')$ is a son of x and $t' := \text{tp}^{\mathcal{J}}(e')$ is as desired, i.e., $t \rightsquigarrow_r t'$ and $\ell(y) = (t', d')$. Condition 3d can be established similarly.

Conversely, assume that there is a $Q \times \Delta^{\mathcal{I}}$ -labelled tree (T, ℓ) that satisfies Conditions 1 to 3. Define an interpretation \mathcal{J}_0 by setting

- $\Delta^{\mathcal{J}_0} = T \setminus \{x \in T \mid \ell_1(x) = q_0\}$;
- $A^{\mathcal{J}_0} = \{x \in \Delta^{\mathcal{J}_0} \mid A \in \ell_1(x)\}$;
- $(x, y) \in r^{\mathcal{J}_0}$ iff y is a son of x , $\ell(x) = (t, d)$, $\ell(y) = (t', d')$, $t \rightsquigarrow_r t'$, and $(d, d') \in r^{\mathcal{I}}$.

The next step is to extend \mathcal{J}_0 to also satisfy existential restrictions $\exists r.C$ with $r \notin \Sigma$. For each $x \in \Delta^{\mathcal{J}_0}$ and $\exists r.C \in \ell_1(x)$ with $r \notin \Sigma$, fix a model $\mathcal{I}_{x, \exists r.C}$ of \mathcal{T} that satisfies C and

every D with $\forall r.D \in \ell_1(x)$ at the root. Such models exist since $\ell_1(x) \in \text{TP}(\mathcal{T})$, thus it is realized in some model of \mathcal{T} . Let $\mathcal{I}_{x_1, \exists r_1.C_1}, \dots, \mathcal{I}_{x_k, \exists r_k.C_k}$ be the chosen models and assume w.l.o.g. that their domains are pairwise disjoint, and also disjoint from $\Delta^{\mathcal{J}_0}$. Now define a new interpretation \mathcal{J} as follows:

- $\Delta^{\mathcal{J}} = \Delta^{\mathcal{J}_0} \cup \bigcup_{1 \leq i \leq k} \Delta^{\mathcal{I}_{x_i, \exists r_i.C_i}}$;
- $A^{\mathcal{J}} = A^{\mathcal{J}_0} \cup \bigcup_{1 \leq i \leq k} A^{\mathcal{I}_{x_i, \exists r_i.C_i}}$;
- $r^{\mathcal{J}} = r^{\mathcal{J}_0} \cup \bigcup_{1 \leq i \leq k} r^{\mathcal{I}_{x_i, \exists r_i.C_i}} \cup \bigcup_{1 \leq i \leq k} (x_i, \rho^{\mathcal{I}_{x_i, \exists r_i.C_i}})$.

By Condition 2, there is a son x_0 of ε in T such that $\ell(x_0) = (t, d_0)$. Using Condition 3, it can be verified that $\{(d, x) \in \Delta^{\mathcal{I}} \times \Delta^{\mathcal{J}_0} \mid \ell_2(x) = d\}$ is a Σ -bisimulation between (\mathcal{I}, d_0) and (\mathcal{J}, x_0) . It thus remains to show that \mathcal{J} is a model of \mathcal{T} , which is an immediate consequence of the following claim and the definition of types for \mathcal{T} .

Claim. For all $C \in \text{cl}(\mathcal{T})$:

- (i) for all $x \in \Delta^{\mathcal{J}_0}$ with $\ell(x) = (t, d)$, $C \in t$ implies $x \in C^{\mathcal{J}}$;
- (ii) for $1 \leq i \leq k$ and all $x \in \Delta^{\mathcal{I}_{x_i, \exists r_i.C_i}}$, $x \in C^{\mathcal{I}_{x_i, \exists r_i.C_i}}$ implies $x \in C^{\mathcal{J}}$.

The proof is by induction on the structure of C . We only do the case $C = \exists r.D$ explicitly. For Point (i), let $x \in \Delta^{\mathcal{J}_0}$ with $\ell(x) = (t, d)$, and $\exists r.D \in t$. First assume $r \in \Sigma$. Then Condition 3d yields a $(d, d') \in r^{\mathcal{I}}$ and $t' \in \text{TP}(\mathcal{T})$ such that $D \in t'$ and $t \rightsquigarrow_r t'$, and a son y of x with $\ell(y) = (t', d')$. By definition of \mathcal{J}_0 , $(x, y) \in r^{\mathcal{J}}$. By IH, $D \in t'$ yields $y \in D^{\mathcal{J}}$, thus $x \in (\exists r.D)^{\mathcal{J}}$. Now assume $r \notin \Sigma$. Then $(x, \rho^{\mathcal{I}_{x, \exists r.D}}) \in r^{\mathcal{J}}$. By IH and choice of $\rho^{\mathcal{I}_{x, \exists r.D}}$, $\rho^{\mathcal{I}_{x, \exists r.D}} \in D^{\mathcal{J}}$ and we are done. For Point (ii), it suffices to apply IH and the semantics. \square

Proof of Theorem 6. We have already proved the modified version of the theorem, where “for all interpretations \mathcal{I} ” is replaced with “for all interpretations \mathcal{I} with finite outdegree”. The “if” direction of the modified version immediately implies the one of the original version. For the “only if” direction of the original version, assume that \mathcal{T}_Σ is a uniform Σ -interpolant of \mathcal{T} . The direction “ \Leftarrow ” of (*) is proved exactly as in the modified version. For “ \Rightarrow ”, take an interpretation \mathcal{I} with $\mathcal{I} \models \mathcal{T}_\Sigma$ and assume to the contrary that there is a $d \in \Delta^{\mathcal{I}}$ such that (\mathcal{I}, d) is not Σ -bisimilar to any pointed model of \mathcal{T} . Let \mathcal{A} be the complement of the automaton $\mathcal{A}_{\mathcal{T}, \Sigma}$ of Theorem 15. By Lemma 28, $(\mathcal{I}, d) \in L(\mathcal{A})$. We can w.l.o.g. assume that \mathcal{I} is a tree interpretation with root d (if it is not, apply unravelling). By considering an accepting run (T, ℓ) of \mathcal{A} on \mathcal{I} and removing unnecessary subtrees synchronously from both \mathcal{I} and (T, ℓ) , it is easy to show that there is a $(\mathcal{J}, d) \in L(\mathcal{A})$ that is still a model of \mathcal{T}_Σ , but of finite outdegree. Since $\mathcal{J} \in L(\mathcal{A})$, \mathcal{J} is not Σ -bisimilar to any model of \mathcal{T} . The existence of such a \mathcal{J} contradicts the modified Theorem 6, which we already proved to hold. \square

Our next aim is to prove Theorem 16, which for convenience we state in expanded form here. **Theorem 16.** Let \mathcal{T} be a TBox, $\Sigma \subseteq \text{sig}(\mathcal{T})$ a signature, and $m \geq 0$. Then there is an APTA $\mathcal{A}_{\mathcal{T}, \Sigma, m} = (Q, \Sigma_N, \Sigma_E, q_0, \delta, \Omega)$ such that $L(\mathcal{A}) \neq \emptyset$ iff there are Σ -tree interpretations (\mathcal{I}_1, d_1) and (\mathcal{I}_2, d_2) such that

1. $\mathcal{I}_1^{\leq m} = \mathcal{I}_2^{\leq m}$;
2. $(\mathcal{I}_1, \rho^{\mathcal{I}_1}) \models \exists_{\Sigma}^{\approx} \mathcal{T}$;
3. $(\mathcal{I}_2, \rho^{\mathcal{I}_2}) \not\models \exists_{\Sigma}^{\approx} \mathcal{T}$;
4. for all successors d of $\rho^{\mathcal{I}_2}$, we have $(\mathcal{I}_2, d) \models \exists_{\Sigma}^{\approx} \mathcal{T}$.

Moreover, $|Q| \in \mathcal{O}(2^{\mathcal{O}(n)} + \log^2 m)$ and $|\Sigma_N|, |\Sigma_E| \in \mathcal{O}(n + \log m)$, where $n = |\mathcal{T}|$.

Proof. Let \mathcal{T} be a TBox, $\Sigma \subseteq \text{sig}(\mathcal{T})$ a signature, and $k \geq 0$. We show how to construct the APTA $\mathcal{A}_{\mathcal{T}, \Sigma, k}$ stipulated in Theorem 16 as an intersection of four automata $\mathcal{A}_1, \dots, \mathcal{A}_4$. Let $k = \lceil \log(m+2) \rceil$. All of the automata \mathcal{A}_i will use the alphabets

$$\begin{aligned} \Sigma_N &= ((\Sigma \cap \mathbf{N}_C) \times \{1, 2\}) \cup \{c_1, \dots, c_k\} \\ \Sigma_E &= (\Sigma \cap \mathbf{N}_R) \times \{1, 2, 12\} \end{aligned}$$

We assume w.l.o.g. that $\Sigma_N \subseteq \mathbf{N}_C$ and $\Sigma_E \subseteq \mathbf{N}_R$. Intuitively, a $\Sigma_N \cup \Sigma_E$ -interpretation \mathcal{I} represents two Σ -interpretations \mathcal{I}_1 and \mathcal{I}_2 where for $i \in \{1, 2\}$, we set $\Delta^{\mathcal{I}_i} = \Delta^{\mathcal{I}}$, $A^{\mathcal{I}_i} = d \in (A, i)^{\mathcal{I}}$ and $r^{\mathcal{I}_i} = (r, i)^{\mathcal{I}} \cup (r, 12)^{\mathcal{I}}$. Thus, edges indexed by “12” represent edges that are shared between the two interpretations. The additional concept names c_1, \dots, c_k are used to implement a counter that counts the depth of elements in \mathcal{I} up to $m+1$, and then stays at $m+1$. As a by-product, the counter ensures that every element has a uniquely defined depth. Formally, a pointed $\Sigma_N \cup \Sigma_E$ -interpretation (\mathcal{I}, d) is called *m-well-counting* if for all $d_0, \dots, d_n \in \Delta^{\mathcal{I}}$ with $d_0 = d$ and $(d_i, d_{i+1}) \in \bigcup_{r \in \Sigma_E} r^{\mathcal{I}}$ for $0 \leq i < n$, the value encoded in binary by the truth of the concept names c_1, \dots, c_k at d_n is $\min\{n, m+1\}$.

Automaton \mathcal{A}_1 ensures Condition 2 of Theorem 16. We construct it by starting with the APTA $\mathcal{A}_{\mathcal{T}, \Sigma}$ from Theorem 15, and then modifying it as follows to run on track 1 of the combined interpretations:

- replace the node alphabet $\Sigma \cap \mathbf{N}_R$ with Σ_N , and the edge alphabet $\Sigma \cap \mathbf{N}_C$ with Σ_E
- for all states q , $\delta(q)$ is obtained from $\delta(q)$ by replacing A with $(A, 1)$, $\neg A$ with $\neg(A, 1)$, $\langle r \rangle q$ with $\langle (r, 1) \rangle q \vee \langle (r, 12) \rangle q$ and every $[r]q$ with $[(r, 1)]q \wedge [(r, 12)]q$

Automaton \mathcal{A}_2 takes care of Condition 3 of Theorem 16. We again start with $\mathcal{A}_{\mathcal{T}, \Sigma}$, first complement it according to Lemma 25 and then modify it as \mathcal{A}_1 , but using track/index 2 instead of track/index 1. Automaton \mathcal{A}_3 addresses Condition 4 of Theorem 16. We start once more with $\mathcal{A}_{\mathcal{T}, \Sigma}$, modify it to run on track 2, add a new initial state q'_0 , and put $\delta(q'_0) = \bigwedge_{r \in (\Sigma \cap \mathbf{N}_R) \times \{2, 12\}} [r]q_0$, with q_0 the original initial state, to start the run of the obtained automaton at every successor of the selected point instead of at the selected point itself. The following lemma states the central property of the APTAs constructed so far.

Claim 1 Let (\mathcal{I}, d) be a pointed $\Sigma_N \cup \Sigma_E$ -interpretation. Then

1. $(\mathcal{I}, d) \in L(\mathcal{A}_1)$ iff $(\mathcal{I}_1, d) \in L(\mathcal{A}_{\mathcal{T}, \Sigma})$ iff $(\mathcal{I}_1, d) \models \exists_{\Sigma}^{\approx} \mathcal{T}$
2. $(\mathcal{I}, d) \in L(\mathcal{A}_2)$ iff $(\mathcal{I}_2, d) \notin L(\mathcal{A}_{\mathcal{T}, \Sigma})$ iff $(\mathcal{I}_2, d) \not\models \exists_{\Sigma}^{\approx} \mathcal{T}$
3. $(\mathcal{I}, d) \in L(\mathcal{A}_3)$ iff $(\mathcal{I}_2(d'), d') \in L(\mathcal{A}_{\mathcal{T}, \Sigma})$ for all successors d' of d iff $(\mathcal{I}_2(d'), d') \models \exists_{\Sigma}^{\approx} \mathcal{T}$ for all successors d' of d .

The purpose of the final automaton \mathcal{A}_4 is to address Condition 1 of Theorem 16. To achieve this, \mathcal{A}_4 also enforces that accepted interpretations are *m-well-counting*.

$$\begin{aligned} Q &= \{q_0, q_1, q_2\} \\ \delta(q_0) &= \neg c_1 \wedge \dots \wedge \neg c_k \wedge q_1 \wedge q_2 \\ \delta(q_1) &= ((c = k + 1) \vee [r](c++)) \wedge \\ &\quad ((c < k + 1) \vee [r](c==)) \wedge \bigwedge_{r \in \Sigma_E} [r]q_1 \\ \delta(q_2) &= \bigwedge_{A \in \Sigma \cap \mathbf{N}_C} ((A, 1) \wedge (A, 2)) \vee (\neg(A, 1) \wedge \neg(A, 2)) \wedge \\ &\quad \bigwedge_{r \in \Sigma \cap \mathbf{N}_R} [(r, 1)]\text{false} \wedge [(r, 2)]\text{false} \\ &\quad \wedge ((c = k + 1) \vee \bigwedge_{r \in \Sigma \cap \mathbf{N}_R} [(r, 12)]q_2) \end{aligned}$$

where $(c < k + 1)$ and $(c = k + 1)$ are the obvious Boolean NNF formulas expressing that the counter c_1, \dots, c_k is smaller and equal to $k + 1$, respectively, $[r](c==)$ is a formula expressing that the counter value does not change when travelling to r -successors, and $[r](c++)$ expresses that the counter is incremented when travelling to r -successors. It is standard to work out the details of these formulas.

Claim 2. Let (\mathcal{I}, d) be a pointed $\Sigma_N \cup \Sigma_E$ -interpretation. Then $(\mathcal{I}, d) \in L(\mathcal{A}_4)$ iff (\mathcal{I}, d) is *k-well-counting* and $(\mathcal{I}_1, d)^{\leq k} = (\mathcal{I}_2, d)^{\leq k}$.¹

By Claims 1 and 2 and Lemma 26, the intersection of $\mathcal{A}_1, \dots, \mathcal{A}_4$ satisfies Conditions 1-4 of Theorem 16. The size bounds stated in Theorem 16 are also satisfied (note the additional states implicit in \mathcal{A}_4). \square

Theorem 17. Given TBoxes \mathcal{T} and \mathcal{T}' , it can be decided in time $2^{p(|\mathcal{T}| \cdot 2^{|\mathcal{T}'|})}$ whether $\mathcal{T} \cup \mathcal{T}'$ is a conservative extension of \mathcal{T} , for some polynomial $p()$.

Proof.(sketch) Let \mathcal{T} and \mathcal{T}' be TBoxes and $\Sigma = \text{sig}(\mathcal{T})$. We construct an APTA \mathcal{A} such that $L(\mathcal{A}) = \emptyset$ iff there are Σ -tree interpretations \mathcal{I} and \mathcal{I}' such that $\mathcal{I} \models \mathcal{T}$ and $\mathcal{I} \not\models \exists_{\Sigma}^{\approx} \mathcal{T}'$. By Theorem 8 and a straightforward unravelling argument, it follows that $\mathcal{T} \cup \mathcal{T}'$ is a conservative extension of \mathcal{T} iff $L(\mathcal{A}) = \emptyset$. We start with taking automata $\mathcal{A}_{\mathcal{T}, \Sigma}$ and $\mathcal{A}_{\mathcal{T}', \Sigma}$, where the latter is constructed according to Theorem 15 and the former is defined as $(Q, \Sigma_N, \Sigma_E, q_0, \delta, \Omega)$, where Q is the

¹The definition of $(\mathcal{I}, d)^{\leq k}$ generalizes from tree interpretations to *k-well-counting* interpretations in the obvious way.

set of all subformulas of $C_{\mathcal{T}}$,

$$\begin{aligned} \Sigma_N &= \Sigma \cap N_C & \Sigma_E &= \Sigma \cap N_R \\ \delta(q_0) &= C_{\mathcal{T}} \wedge \bigwedge_{r \in \Sigma \cap N_R} [r]q_0 \\ \delta(A) &= A & \delta(\neg A) &= \neg A \\ \delta(C \sqcap D) &= C \wedge D & \delta(C \sqcup D) &= C \vee D \\ \delta(\exists r.C) &= \langle r \rangle C & \delta(\forall r.C) &= [r]C \end{aligned}$$

Set $\Omega(q) = 0$ for all $q \in Q$. We can then use the constructions from Lemma 25 to obtain a final automaton \mathcal{A} such that $L(\mathcal{A}) = L(\mathcal{A}_{\mathcal{T}, \Sigma}) \cap \bar{L}(\mathcal{A}_{\mathcal{T}', \Sigma})$. It can be verified that the automaton is as required, and that Theorem 27 yields the bounds stated in Theorem 17. \square

Theorem 18. It is 2-EXPTIME-complete to decide for a given a TBox \mathcal{T} and signature $\Sigma \subseteq \text{sig}(\mathcal{T})$, whether there exists a uniform Σ -interpolant of \mathcal{T} .

Proof. It remains to prove the lower bound. To this end, we reduce deciding conservative extensions to deciding the existence of uniform interpolants. Assume $\mathcal{T} \subseteq \mathcal{T}'$ are given. We may assume that $\mathcal{T}' = \mathcal{T} \cup \{\top \sqsubseteq C\}$ and that \mathcal{T}' is satisfiable and $\mathcal{T} \not\models \mathcal{T}'$. Consider the TBox

$$\mathcal{T}_0 = \mathcal{T} \cup \{\neg C \sqsubseteq A, A \sqsubseteq \exists r.A\} \cup \{\exists s.A \sqsubseteq A \mid s \in \text{sig}(\mathcal{T}')\}$$

where A is a fresh concept name and r a fresh role name.

Claim. \mathcal{T}' is a conservative extension of \mathcal{T} iff there exists a uniform Σ -interpolant of \mathcal{T}_0 for $\Sigma = \text{sig}(\mathcal{T}) \cup \{r\}$.

Assume first that \mathcal{T}' is a conservative extension of \mathcal{T} . We show that \mathcal{T} is a uniform Σ -interpolant of \mathcal{T}_0 . By Theorem 6, it is sufficient to show the following for all $\mathcal{I}: \mathcal{I} \models \mathcal{T}$ iff for all $d \in \Delta^{\mathcal{I}}: (\mathcal{I}, d) \models \exists \Sigma. \mathcal{T}_0$. The direction from right to left is trivial. Assume now that $\mathcal{I} \models \mathcal{T}$ and fix a $d \in \Delta^{\mathcal{I}}$. By Theorem 8, $(\mathcal{I}, d) \models \exists_{\text{sig}(\mathcal{T})} \mathcal{T}'$ and thus there is a pointed model (\mathcal{J}, e) of \mathcal{T}' such that $(\mathcal{I}, d) \sim_{\text{sig}(\mathcal{T})} (\mathcal{J}, e)$. As $\top \sqsubseteq C \in \mathcal{T}'$, we have $C^{\mathcal{J}} = \Delta^{\mathcal{J}}$. Moreover, since A and r do not occur in \mathcal{T} and \mathcal{T}' , we may assume that $A^{\mathcal{J}} = \emptyset$. But from $C^{\mathcal{J}} = \Delta^{\mathcal{J}}$ and $A^{\mathcal{J}} = \emptyset$ we obtain $\mathcal{J} \models \mathcal{T}_0$ and so $(\mathcal{I}, d) \models \exists \Sigma. \mathcal{T}_0$ as required.

Conversely, assume that \mathcal{T}' is not a conservative extension of \mathcal{T} . Let C_0 be a $\text{sig}(\mathcal{T})$ -concept that is satisfiable w.r.t. \mathcal{T} but not w.r.t. \mathcal{T}' . Then $\mathcal{T}_0 \models C_0 \sqsubseteq A$. To show this, let \mathcal{I} be a tree interpretation satisfying \mathcal{T}_0 with $\rho^{\mathcal{I}} \in C_0^{\mathcal{I}}$. Let \mathcal{I}_0 be the restriction of \mathcal{I} to all $d \in \Delta^{\mathcal{I}}$ that are reachable from $\rho^{\mathcal{I}}$ with paths using roles from $\text{sig}(\mathcal{T}')$ only. Then \mathcal{I}_0 is a model of \mathcal{T} and not a model of \mathcal{T}' and so $C^{\mathcal{I}_0} \neq \Delta^{\mathcal{I}_0}$. From $\neg C \sqsubseteq A \in \mathcal{T}_0$ we obtain $A^{\mathcal{I}_0} \neq \emptyset$. Thus, from $\exists s.A \sqsubseteq A \in \mathcal{T}_0$ for all $s \in \text{sig}(\mathcal{T}')$, we obtain $\rho^{\mathcal{I}} \in A^{\mathcal{I}_0} \subseteq A^{\mathcal{I}}$, as required.

From $A \sqsubseteq \exists r.A \in \mathcal{T}_0$, we obtain $\mathcal{T}_0 \models C_0 \sqsubseteq \exists r. \dots \exists r. \top$ for arbitrary long sequences $\exists r. \dots \exists r$. Intuitively, this cannot be axiomatized with a uniform interpolant not using A . To prove this in a formal way, we apply Theorem 9 and prove that $(*_m)$ holds for all $m > 0$.

Assume $m > 0$ is given. Let \mathcal{J}_0 be a tree model satisfying \mathcal{T} with $C_0^{\mathcal{J}_0} \in \rho^{\mathcal{J}_0}$ and in which each $d \in \Delta^{\mathcal{J}_0}$ with the exception of $\rho^{\mathcal{J}_0}$ has an r -successor. Let \mathcal{J}_i be tree models satisfying \mathcal{T}' with $r^{\mathcal{J}_i} = \emptyset, i > 0$. We may assume that the $\Delta^{\mathcal{J}_i}$ are mutually disjoint. Define $\mathcal{I}_1 = (\Delta^{\mathcal{I}_1}, \cdot^{\mathcal{I}_1})$ by setting

- $\Delta^{\mathcal{I}_1} = \bigcup_{i \geq 0} \Delta^{\mathcal{J}_i}$,
- $s^{\mathcal{I}_1} = \bigcup_{i \geq 0} s^{\mathcal{J}_i}$ and $B^{\mathcal{I}_1} = \bigcup_{i \geq 0} B^{\mathcal{J}_i}$ for all $s, B \in \text{sig}(\mathcal{T}')$;
- $r^{\mathcal{I}_1} = r^{\mathcal{J}_0} \cup \{(\rho^{\mathcal{J}_i}, \rho^{\mathcal{J}_{i+1}}) \mid i \geq 0\}$.

Note that $\rho^{\mathcal{I}_1} = \rho^{\mathcal{J}_0}$. Then $(\mathcal{I}_1, \rho^{\mathcal{I}_1}) \models \exists \Sigma. \mathcal{T}_0$ because the extension of \mathcal{I}_1 defined by setting $A^{\mathcal{I}_1} = \Delta^{\mathcal{J}_0} \cup \{\rho^{\mathcal{J}_i} \mid i > 0\}$ is a model of \mathcal{T}_0 .

Let \mathcal{I}_2 be the restriction of \mathcal{I}_1 to $\bigcup_{0 \leq i \leq m} \Delta^{\mathcal{J}_i}$. Then $(\mathcal{I}_2, \rho^{\mathcal{I}_2}) \not\models \exists \Sigma. \mathcal{T}_0$: for any interpretation \mathcal{I}'_2 with $(\mathcal{I}_2, \rho^{\mathcal{I}_2}) \sim_{\Sigma} (\mathcal{I}'_2, \rho^{\mathcal{I}'_2})$ we have $\rho^{\mathcal{I}'_2} \in C_0^{\mathcal{I}'_2}$. If \mathcal{I}'_2 is a model of \mathcal{T}_0 , then $\rho^{\mathcal{I}'_2} \in A^{\mathcal{I}'_2}$ and so there exist d_0, d_1, \dots such that $d_0 = \rho^{\mathcal{I}'_2}$ and $(d_i, d_{i+1}) \in r^{\mathcal{I}'_2}$ for $i \geq 0$. But then there exists such a sequence in \mathcal{I}_2 starting at $\rho^{\mathcal{I}_2}$. As such a sequence does not exist, we have derived a contradiction.

On the other hand, for all sons d of $\rho^{\mathcal{I}_2}$, we have $(\mathcal{I}_2, d) \models \exists \Sigma. \mathcal{T}_0$: for $d = \rho^{\mathcal{J}_1}$ this is witnessed by the interpretation obtained from \mathcal{I}_2 by interpreting A as the empty set. For all $d \in \Delta^{\mathcal{J}_0}$ this is witnessed by the interpretation obtained from \mathcal{I}_2 by interpreting A as the whole domain.

It follows that \mathcal{I}_1 and \mathcal{I}_2 satisfy the condition $(*_m)$ from Theorem 9. \square

10 Proofs for Section 5

Theorem 20. Let $\mathcal{T} = \{\top \sqsubseteq C_{\mathcal{T}}\}$ and assume that \mathcal{T} has a uniform Σ -interpolant. Let R denote the set of role names in \mathcal{T} , $m = 2^{2^{|\mathcal{C}_{\mathcal{T}}|+1}} + 2^{|\mathcal{C}_{\mathcal{T}}|} + 2$ and let C be a Σ -concept uniform interpolant of $\forall R^{\leq m}. C_{\mathcal{T}}$ w.r.t. Σ . Then $\mathcal{T}' = \{\top \sqsubseteq C\}$ is a uniform Σ -interpolant of \mathcal{T} .

Proof. Recall that $M_{\mathcal{T}} = 2^{2^{|\mathcal{C}_{\mathcal{T}}|}}$. By Theorem 12, $\mathcal{T}_{\Sigma, M_{\mathcal{T}}^2+1}$ is a uniform Σ -interpolant of \mathcal{T} . We may assume that $\mathcal{T}_{\Sigma, M_{\mathcal{T}}^2+1} = \{\top \sqsubseteq F\}$ for a Σ -concept F with $\text{rd}(F) \leq M_{\mathcal{T}}^2 + 1$. We show

$$\emptyset \models \forall R^{\leq m}. C_{\mathcal{T}} \sqsubseteq F$$

We provide a sketch only since the argument is similar to the standard reduction of “global consequence” to “local consequence” in modal logic. Suppose this is not the case. Let \mathcal{I} be a tree interpretation with $\rho^{\mathcal{I}} \in (\forall R^{\leq m}. C_{\mathcal{T}})^{\mathcal{I}}$ and $\rho^{\mathcal{I}} \notin F^{\mathcal{I}}$. Let W be the set of $d \in \Delta^{\mathcal{I}}$ that are of depth $2^{2^{|\mathcal{C}_{\mathcal{T}}|+1}}$. For any path of length $2^{|\mathcal{C}_{\mathcal{T}}|} + 1$ starting at some $d \in W$, there exist at least two points on that path, say d_1 and d_2 , such that

$$\{E \in \text{sub}(C_{\mathcal{T}}) \mid d_1 \in E^{\mathcal{I}}\} = \{E \in \text{sub}(C_{\mathcal{T}}) \mid d_2 \in E^{\mathcal{I}}\}.$$

We remove the subtree $\mathcal{I}(d_2)$ from \mathcal{I} and add the pair (d', d_1) to $r^{\mathcal{I}}$ for the unique predecessor d' of d_2 with $(d', d_2) \in r^{\mathcal{I}}$ for some role r . This modification is repeated until a (non-tree!) interpretation \mathcal{I}' is reached in which all points are reachable from $\rho^{\mathcal{I}'}$ by a path of length bounded by $m = 2^{2^{|\mathcal{C}_{\mathcal{T}}|+1}} + 2^{|\mathcal{C}_{\mathcal{T}}|} + 2$. Since $\text{rd}(F) \leq M_{\mathcal{T}}^2 + 1$ and \mathcal{I} has not changed for points of depth not exceeding $M_{\mathcal{T}}^2 + 1$, we still have $\rho^{\mathcal{I}'} \notin F^{\mathcal{I}'}$. By construction, \mathcal{I}' is a model of $\mathcal{T} = \{\top \sqsubseteq C_{\mathcal{T}}\}$. Thus, we have obtained a contradiction

to the assumption that $\{\top \sqsubseteq F\}$ is a uniform Σ -interpolant of \mathcal{T} .

From $\emptyset \models \forall R^{\leq m}. C_{\mathcal{T}} \sqsubseteq F$ we obtain $\emptyset \models C \sqsubseteq F$ for the Σ -concept uniform interpolant C . Thus $\{\top \sqsubseteq C\} \models \top \sqsubseteq F$ and so $\{\top \sqsubseteq C\}$ is a uniform Σ -interpolant of \mathcal{T} . \square

Theorem 21. There exists a signature Σ and a family of TBoxes $(\mathcal{T}_n)_{n>0}$ such that, for all $n > 0$,

- (i) $|\mathcal{T}_n| \in \mathcal{O}(n^2)$ and
- (ii) every uniform Σ -interpolant $\{\top \sqsubseteq C_{\mathcal{T}}\}$ for \mathcal{T}_n is of size at least $2^{(2^n \cdot 2^{2^n})-2}$.

To prove Theorem 21 in an economic way, we reuse some techniques and result from [6]. We first need a bit of terminology. If \mathcal{T} and \mathcal{T}' are TBoxes and \mathcal{T}' is not a conservative extension of \mathcal{T} , then there is a Σ -concept C such that C is satisfiable relative \mathcal{T} , but not relative to \mathcal{T}' ; such a concept C is a *witness concept* for non-conservativity of the extension of \mathcal{T} with \mathcal{T}' . One main result of [6] is as follows.

Theorem 29 ([6]). *There are families of TBoxes $(\mathcal{T}_n)_{n>0}$ and $(\mathcal{T}'_n)_{n>0}$ such that, for all $n > 0$,*

- (i) $\mathcal{T}_n \cup \mathcal{T}'_n$ is not a conservative extension of \mathcal{T}_n ,
- (i) $|\mathcal{T}_n| \in \mathcal{O}(n^2)$, $|\mathcal{T}'_n| \in \mathcal{O}(n^2)$, and
- (ii) every witness concept for non-conservativity of the extension of \mathcal{T}_n with \mathcal{T}'_n is of size at least $2^{(2^n \cdot 2^{2^n})-1}$.

To transfer Theorem 29 from witness concepts to uniform interpolants, we need to introduce some technicalities from its proof. For the remainder of this section, fix a signature $\Sigma = \{A, B, r, s\}$. Let \mathcal{I} be an interpretation and $d \in \Delta^{\mathcal{I}}$. A *path* starting at d is a sequence d_1, \dots, d_k with $d_1 = d$ and $(d_i, d_{i+1}) \in r^{\mathcal{I}} \cup s^{\mathcal{I}}$, for $1 \leq i \leq k$. In [6], \mathcal{I} is called *strongly n -violating* iff there exists an $x \in A^{\mathcal{I}}$ such that the following two properties are satisfied, where $m = (2^n \cdot 2^{2^n})$:

- (P1) for all paths x_1, \dots, x_k in \mathcal{I} with $k \leq m$ starting at x , the X values of x_1, \dots, x_k describe the first k bits of a 2^n -bit counter counting from 0 to $2^{2^n} - 1$.
- (P2) there exist elements $x_w \in \Delta^{\mathcal{I}}$, for all $w \in \{r, s\}^*$ of length at most $m - 1$, such that the following are true:
 - (a) $x_\varepsilon = x$;
 - (b) $(x_w, x_{w'}) \in r^{\mathcal{I}}$ if $w' = w \cdot r$, and $(x_w, x_{w'}) \in s^{\mathcal{I}}$ if $w' = w \cdot s$;
 - (c) $x_w \notin B^{\mathcal{I}}$ if w is of length $m - 1$.

Define a Σ -TBox

$$\mathcal{T}_n^- = \{\top \sqsubseteq \forall r. \neg A \cap \forall s. \neg A \\ A \sqsubseteq \neg X \cap \bigcap_{i < 2^n} \forall (r \cup s)^i. \neg X\}$$

The following result of [6] underlies the proof of Theorem 29.

Lemma 30 ([6]). *There exist families of TBoxes $(\mathcal{T}_n)_{n>0}$ and $(\mathcal{T}'_n)_{n>0}$ such that, for all $n > 0$,*

- (i) $|\mathcal{T}_n| \in \mathcal{O}(n^2)$, $|\mathcal{T}'_n| \in \mathcal{O}(n^2)$;
- (ii) a model of \mathcal{T}_n that is strongly n -violating cannot be extended to a model of \mathcal{T}'_n ;

$$K_1 = \bigcap_{i < m, \text{bit}_i(2^{2^n})=1} \forall \{r, s\}^i. X \cap \quad (1)$$

$$\bigcap_{i < m, \text{bit}_i(2^{2^n})=0} \forall \{r, s\}^i. \neg X \quad (2)$$

$$K_2^{(0)} = \neg B \quad (3)$$

$$K_2^{(i+1)} = \exists r. K_2^{(i)} \cap \exists s. K_2^{(i)} \quad (4)$$

Figure 1: Definition of the concepts K_1 and $K_2^{(m)}$.

(iii) a tree model of \mathcal{T}_n that is not strongly n -violating can be extended to a model of \mathcal{T}'_n ;

(iv) every model of \mathcal{T}_n^- can be extended to a model of \mathcal{T}_n .

The TBoxes \mathcal{T}_n and \mathcal{T}'_n from Lemma 29 are formulated in extensions of the signature Σ , more precisely we have $\Sigma = \text{sig}(\mathcal{T}_n) \cap \text{sig}(\mathcal{T}'_n)$. Thus, the phrase ‘extended to a model of’ refers to interpreting those symbols that do not occur in the original TBox.

To establish Theorem 21, we consider the uniform Σ -interpolants of the TBoxes $\mathcal{T}_n \cup \mathcal{T}'_n$. Let

$$\mathcal{T}_{\Sigma, n} = \mathcal{T}_n^- \cup \{A \sqsubseteq \neg K_1 \sqcup \neg K_2^{(m)}\}$$

where the concepts K_1 and $K_2^{(m)}$ are shown in Figure 1 and $m = 2^n \cdot 2^{2^n}$. In the figure, we use $\text{bit}_i(m)$ to denote the i -th bit of the string obtained by concatenating all values of a binary counter that counts up to m (lowest bit first, and with every counter value padded to $\log(m)$ bits using trailing zeros). Note that $\forall \{r \cup s\}^i. C$ is an abbreviation for

$\bigcap_{r_1 \dots r_i \in \{r, s\}^i} \forall r_1. \dots \forall r_i. C$. We show that $\mathcal{T}_{\Sigma, n}$ is a uniform

Σ -interpolant of $\mathcal{T}_n \cup \mathcal{T}'_n$, and that it is essentially of minimal size. It is not hard to see that the models of $\mathcal{T}_{\Sigma, n}$ are precisely those interpretations that are not strongly n -violating.

Lemma 31. *For all $n \geq 0$,*

1. $\mathcal{T}_{\Sigma, n}$ is a uniform Σ -interpolant of $\mathcal{T}_n \cup \mathcal{T}'_n$;
2. every uniform Σ -interpolant $\{\top \sqsubseteq C_{\mathcal{T}}\}$ for $\mathcal{T}_n \cup \mathcal{T}'_n$ is of size at least $2^{(2^n \cdot 2^{2^n})-1}$.

Proof. For Point 1, let $C \sqsubseteq D$ be a Σ -inclusion and assume first that $\mathcal{T}_{\Sigma, n} \not\models C \sqsubseteq D$, i.e., there is a model \mathcal{I} of $\mathcal{T}_{\Sigma, n}$ and a $d \in (C \cap \neg D)^{\mathcal{I}}$. Since $\mathcal{T}_n^- \subseteq \mathcal{T}_{\Sigma, n}$ and by Point (iv) of Lemma 30, \mathcal{I} can be extended to a model \mathcal{I}' of \mathcal{T}_n . Let \mathcal{J} be the unravelling of \mathcal{I}' into a tree with root d . Obviously, \mathcal{J} is still a model of \mathcal{T}_n and $d \in (C \cap \neg D)^{\mathcal{J}}$. Since $\mathcal{I} \models \mathcal{T}_{\Sigma, n}$, \mathcal{I} is not strongly n -violating, and thus the same holds for \mathcal{J} . By Point (iii) of Lemma 30, \mathcal{J} can be extended to a model \mathcal{J}' of \mathcal{T}'_n , and we have $d \in (C \cap \neg D)^{\mathcal{J}'}$, thus $\mathcal{T}_n \cup \mathcal{T}'_n \not\models C \sqsubseteq D$.

Conversely, let $\mathcal{T}_n \cup \mathcal{T}'_n \not\models C \sqsubseteq D$. Then there is a model \mathcal{I} of $\mathcal{T}_n \cup \mathcal{T}'_n$ and a $d \in (C \cap \neg D)^{\mathcal{I}}$. By (ii), \mathcal{I} is not strongly n -violating, thus it is a model of $\mathcal{T}_{\Sigma, n}$ and we get $\mathcal{T}_{\Sigma, n} \not\models C \sqsubseteq D$.

For Point 2, assume that there is a uniform Σ -interpolant $\{\top \sqsubseteq C_{\mathcal{T}}\}$ for $\mathcal{T}_n \cup \mathcal{T}'_n$ that is of size strictly smaller than $2^{(2^n \cdot 2^{2^n})-2}$. Then the size of $\neg C_{\mathcal{T}}$ is strictly smaller than

$2^{(2^n \cdot 2^{2^n}) - 1}$. By Theorem 29, to obtain a contradiction it thus suffices to show that $\neg C_{\mathcal{T}}$ is a witness concept for non-conservativity of the extension of \mathcal{T}_n with \mathcal{T}'_n . First, since $\mathcal{T}_n \cup \mathcal{T}'_n \models \top \sqsubseteq C_{\mathcal{T}}$, $\neg C_{\mathcal{T}}$ is unsatisfiable relative to $\mathcal{T}_n \cup \mathcal{T}'_n$. And second, there clearly is a model \mathcal{I} of \mathcal{T}_n^- that is not a model of $\mathcal{T}_{\Sigma, n}$. Since $\{\top \sqsubseteq C_{\mathcal{T}}\}$ and $\mathcal{T}_{\Sigma, n}$ are both uniform Σ -interpolants of $\mathcal{T}_n \cup \mathcal{T}'_n$, they are equivalent and thus there is a $d \in \neg C_{\mathcal{T}}$. By Point (iv) of Lemma 30, $\neg C_{\mathcal{T}}$ is satisfiable relative to \mathcal{T}_n . \square