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 [Grau and Motik, 2010]. 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 [Reiter and Lin, 1994]. In propositional logic, forgetting is also known as *variable elimination* [Lang *et al.*, 2003]. 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 und uniform interpolation receive increased interest also in a DL context [Eiter *et al.*, 2006; Wang *et al.*, 2010; 2009b; 2008; Kontchakov *et al.*, 2010; Konev *et al.*, 2009]. 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 [Kontchakov *et al.*, 2010; Konev *et al.*, 2009] and practical experiments have confirmed the usefulness and feasibility of their computation [Konev *et al.*, 2009]. 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 [ten Cate *et al.*, 2006; Wang *et al.*, 2009b], which is also what the literature on uniform interpolants in modal logic is about [Visser, 1996; Herzig and Mengin, 2008]. 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) [Ghilardi *et al.*, 2006]. A scheme for approximating (existing or non-existing) interpolants of $\mathcal{ALC}$-TBoxes was devised in [Wang *et al.*, 2008]. In [Wang *et al.*, 2010], 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 [Wang *et al.*, 2010] 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 [Ghilardi *et al.*, 2006]. Moreover, we use our approach to reprove the 2-EXPTIME upper bound for deciding conservative extensions from [Ghilardi *et al.*, 2006], in an alternative and argueably 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 $\mathsf{N_C}$ and $\mathsf{N_R}$ 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 $A \in \mathsf{N_C}$ and $r \in \mathsf{N_R}$. The concept constructors $\bot$, $\sqcup$, and $\forall r.C$ are defined as abbreviations: $\bot$ 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 $A \in \mathsf{N_C}$ a set $A^\mathcal{I} \subseteq \Delta^\mathcal{I}$, and to each $r \in \mathsf{N_R}$ a relation $r^\mathcal{I} \subseteq \Delta^\mathcal{I} \times \Delta^\mathcal{I}$. The interpretation function is inductively extended to concepts as follows:

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

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 $\Sigma \subseteq \mathsf{N_C} \cup \mathsf{N_R}$ of concept and role names is called a *signature*. The signature $\mathsf{sig}(C)$ of a concept C is the set of concept and role names occurring in C, and likewise for the signature $\mathsf{sig}(C \sqsubseteq D)$ of an inclusion $C \sqsubseteq D$ and $\mathsf{sig}(\mathcal{T})$ of a TBox $\mathcal{T}$. A Σ-*TBox* is a TBox with $\mathsf{sig}(\mathcal{T}) \subseteq \Sigma$, 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 $\Sigma = \mathsf{sig}(\mathcal{T})$.
- $\mathcal{T}$ a *uniform* Σ-*interpolant* of $\mathcal{T}'$ if $\mathsf{sig}(\mathcal{T}) \subseteq \Sigma \subseteq \mathsf{sig}(\mathcal{T}')$ 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 [Wang *et al.*, 2010] 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 $\mathsf{sig}(\mathcal{T}) \setminus \Sigma$-interpolant of $\mathcal{T}$.

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

(1) $\mathsf{Uni} \sqsubseteq \exists \mathsf{has\_st.Undergrad} \sqcap \exists \mathsf{has\_st.Grad}$

(2) $\mathsf{Uni} \sqcap \mathsf{Undergrad} \sqsubseteq \bot$ (3) $\mathsf{Uni} \sqcap \mathsf{Grad} \sqsubseteq \bot$

(4) $\mathsf{Undergrad} \sqcap \mathsf{Grad} \sqsubseteq \bot$.

Then the TBox that consists of (2) and

$\mathsf{Uni} \sqsubseteq \exists \mathsf{has\_st.Undergrad} \sqcap \exists \mathsf{has\_st.}(\neg\mathsf{Undergrad} \sqcap \neg\mathsf{Uni})$

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

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 $\{B\}$.

(i) 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 $\mathcal{T}_1$.

(ii) 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 $\mathcal{T}_2$.

(iii) 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 $\mathcal{T}_3$.

(iv) 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 [Konev *et al.*, 2009].

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 $A \in \Sigma \cap \mathsf{N_C}$;
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 $r \in \Sigma \cap \mathsf{N_R}$;
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 $r \in \Sigma \cap \mathsf{N_C}$.

$(\mathcal{I}_1, d_1)$ 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 [Goranko and Otto, 2007]. Say that $(\mathcal{I}_1, d_1)$ and $(\mathcal{I}_2, d_2)$ are *$\mathcal{ALC}_\Sigma$-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 $\{d' \mid (d, d') \in \bigcup_{r \in \mathsf{N_R}} r^\mathcal{I}\}$ 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 $(\mathcal{I}, d) \models \exists^\sim_\Sigma.\mathcal{T}$ 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 [French, 2006].

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

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

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 $(\Delta^\mathcal{I}, \bigcup_{r \in \mathsf{N_R}} r^\mathcal{I})$ is a (possibly infinite) tree and $r^\mathcal{I} \cap s^\mathcal{I} = \emptyset$ for all distinct $r, s \in \mathsf{N_R}$.

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 $\forall d \in \Delta^\mathcal{I}$: $(\mathcal{I}, d) \models \exists^\sim_{\Sigma_1}.\mathcal{T}_1$.

The first 'iff' relies on the fact that unraveling an interpretation into a tree interpretation preserves bisimularity, 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 $\mathcal{I}$, $\mathcal{I} \models \mathcal{T} \Rightarrow$ for all $d \in \Delta^\mathcal{I}$, $(\mathcal{I}, d) \models \exists^\sim_\Sigma.\mathcal{T}'$ where $\Sigma = \mathsf{sig}(\mathcal{T})$.*

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 $\mathcal{T}^\infty_\Sigma$ 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* $\mathsf{rd}(C)$ 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 $\mathcal{C}^m_f(\Sigma)$ of Σ-concepts D with $\mathsf{rd}(D) \leq m$ such that every Σ-concept C with $\mathsf{rd}(C) \leq m$ is equivalent to some $D \in \mathcal{C}^m_f(\Sigma)$. Let

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

Clearly, $\mathcal{T}^\infty_\Sigma$ 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 [Wang *et al.*, 2008]):

(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 $(\Delta^\mathcal{I}, \bigcup_{r \in \mathsf{N_R}} r^\mathcal{I})$. 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 $(\Delta^\mathcal{I}, \bigcup_{r \in \mathsf{N_R}} r^\mathcal{I})$.

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, $\Sigma \subseteq \mathsf{sig}(\mathcal{T})$, 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. *$(\mathcal{I}_1, \rho^{\mathcal{I}_1}) \models \exists^\sim_\Sigma.\mathcal{T}$;*
3. *$(\mathcal{I}_2, \rho^{\mathcal{I}_2}) \not\models \exists^\sim_\Sigma.\mathcal{T}$;*
4. *For all successors d of $\rho^{\mathcal{I}_2}$: $(\mathcal{I}_2, d) \models \exists^\sim_\Sigma.\mathcal{T}$.*

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 $\mathcal{I}_1 = (\{0, 1, \ldots\}, A^{\mathcal{I}_1}, r^{\mathcal{I}_1})$, 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 $\{0, \ldots, m\}$. Then (1) $\mathcal{I}_1^{\leq m} = \mathcal{I}_2^{\leq m}$; (2) $(\mathcal{I}_1, 0) \models \exists^{\sim}_{\overline{\Sigma}_3} \mathcal{T}_3$ as the expansion of $\mathcal{I}_1$ by $B^{\mathcal{I}_1} = \{0, 1, \ldots\}$ is a model of $\mathcal{T}_3$; (3) $(\mathcal{I}_2, 0) \not\models \exists^{\sim}_{\overline{\Sigma}_3} \mathcal{T}_3$ as there is no infinite r-sequence in $\mathcal{I}_2$ starting at 0; and (4) $(\mathcal{I}_2, 1) \models \exists^{\sim}_{\overline{\Sigma}_3} \mathcal{T}_3$ as the restriction of $\mathcal{I}_2$ to $\{1, \ldots, m\}$ 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 $\mathsf{conc}(\mathcal{T}) = \{C, D \mid C \sqsubseteq D \in \mathcal{T}\}$ 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 $\mathsf{sig}(\mathcal{T}) \setminus \Sigma$ 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 $m = \max\{\mathsf{rd}(C) \mid C \in \mathsf{conc}(\mathcal{T})\}$.

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

It suffices to show that $(*_{M_{\mathcal{T}}^2+1})$ implies $(*_m)$ for all $m \geq M_{\mathcal{T}}^2 + 1$. The proof idea is as follows. Denote by $\mathsf{cl}(\mathcal{T})$ the closure under single negation and subconcepts of $\mathsf{conc}(\mathcal{T})$. The *type* of some $d \in \Delta^{\mathcal{I}}$ in an interpretation $\mathcal{I}$ is

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

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 $(*_{M_{\mathcal{T}}^2+1})$ into a witness $\mathcal{I}_1', \mathcal{I}_2'$ for $(*_m)$ when $m \geq M_{\mathcal{T}}^2 + 1$. 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* $\mathsf{Ext}^{\mathcal{I}}(d)$, defined as

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

and capturing all ways in which the restiction of $\mathsf{tp}^{\mathcal{I}}(d)$ 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 [Ghilardi *et al.*, 2006], this yields decidablity 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 [Ghilardi *et al.*, 2006].

We use amorphous alternating parity tree automata in the style of Wilke [Wilke, 2001], 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*, $\Sigma_N \subseteq \mathsf{N_C}$ is the finite *node alphabet*, $\Sigma_E \subseteq \mathsf{N_R}$ is the finite *edge alphabet*, $q_0 \in Q$ is the *initial state*, $\delta : Q \to \mathsf{mov}(\mathcal{A})$, is the transition function with $\mathsf{mov}(\mathcal{A}) = \{\mathsf{true}, \mathsf{false}, A, \neg A, q, q \wedge q', q \vee q', \langle r \rangle q, [r]q \mid A \in \Sigma_N, q, q' \in Q, r \in \Sigma_E\}$ the set of *moves* of the automaton, and $\Omega : Q \to \mathbb{N}$ 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 $\ell : T \to \Sigma$ 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 $\ell(x) = (q, d)$:

- $\delta(q) \neq \mathsf{false}$;
- 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 [Wilke, 2001].

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 $\Sigma \subseteq \mathsf{sig}(\mathcal{T})$ a signature. Then there exists an APTA $\mathcal{A}_{\mathcal{T},\Sigma} = (Q, \Sigma \cap \mathsf{N_C}, \Sigma \cap \mathsf{N_R}, q_0, \delta, \Omega)$ 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 $(\mathcal{I}, d) \models \exists^{\sim}_{\Sigma}.\mathcal{T}$. $\mathcal{A}_{\mathcal{T},\Sigma}$ 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 [D'Agostino and Hollenberg, 1998], 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 [Baader *et al.*, 2003]. Recall the notion of a type introduced in Section 3. We use $\mathsf{TP}(\mathcal{T})$ to denote the set of all types realized in some model of $\mathcal{T}$, i.e., $\mathsf{TP}(\mathcal{T}) = \{\mathsf{tp}^{\mathcal{I}}(d) \mid \mathcal{I} \text{ model of } \mathcal{T}, d \in \Delta^{\mathcal{I}}\}$. Note that $\mathsf{TP}(\mathcal{T})$ can be computed in time exponential in the size of $\mathcal{T}$ since concept satisfiability w.r.t. TBoxes is EXPTIME-complete in $\mathcal{ALC}$ [Baader *et al.*, 2003]. Given $t, t' \in \mathsf{TP}(\mathcal{T})$ and $r \in \Sigma$, we write $t \rightsquigarrow_r t'$ if $C \in t'$ implies $\exists r.C \in t$ for all $\exists r.C \in \mathsf{cl}(\mathcal{T})$. Now define the automaton $\mathcal{A}_{\mathcal{T},\Sigma} := (Q, \Sigma_N, \Sigma_E, q_0, \delta, \Omega)$, where

$$Q = \mathsf{TP}(\mathcal{T}) \uplus \{q_0\} \quad \Sigma_N = \Sigma \cap \mathsf{N_C} \quad \Sigma_E = \Sigma \cap \mathsf{N_R}$$

$$\delta(q_0) = \bigvee \mathsf{TP}(\mathcal{T})$$

$$\begin{aligned}\delta(t) = & \bigwedge_{A \in t \cap \mathsf{N_C} \cap \Sigma} A \wedge \bigwedge_{A \in (\mathsf{N_C} \cap \Sigma) \setminus t} \neg A \\ & \wedge \bigwedge_{r \in \Sigma \cap \mathsf{N_R}} [r] \bigvee \{t' \in \mathsf{TP}(\mathcal{T}) \mid t \rightsquigarrow_r t'\} \\ & \wedge \bigwedge_{\exists r.C \in t, r \in \Sigma} \langle r \rangle \bigvee \{t' \in \mathsf{TP}(\mathcal{T}) \mid C \in t \wedge t \rightsquigarrow_r t'\}\end{aligned}$$

$$\Omega(q) = 0 \text{ for all } q \in Q$$

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, $\Sigma \subseteq \mathsf{sig}(\mathcal{T})$ 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 [Ghilardi *et al.*, 2006]. 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 [Ghilardi *et al.*, 2006], 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 $\mathsf{sig}(\mathcal{T}) \cup \{r\}$-interpolant of

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

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 $\Sigma \subseteq \mathsf{sig}(\mathcal{T})$, 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 $\mathcal{T}_{\Sigma, M^2_{\mathcal{T}}+1}$ and establish a matching upper bound on their size. Let C be a concept and $\Sigma \subseteq \mathsf{sig}(C)$ a signature. A concept C' is called a *concept uniform Σ-interpolant of* C if $\mathsf{sig}(C) \subseteq \Sigma$, $\emptyset \models C \sqsubseteq C'$, and $\emptyset \models C' \sqsubseteq D$ for every concept D such that $\mathsf{sig}(D) \subseteq \Sigma$ and $\emptyset \models C \sqsubseteq D$. The following result is proved in [ten Cate *et al.*, 2006].

Theorem 19. *For every concept C and signature $\Sigma \subseteq \mathsf{sig}(C)$ 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 [Wang *et al.*, 2010], 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

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

It can be shown using Theorem 12 that for $m = 2^{2^{|C_\mathcal{T}|+1}} + 2^{|C_\mathcal{T}|} + 2$ 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 $\mathsf{rd}(C) \leq \mathsf{rd}(\forall R^{\leq m}.C_\mathcal{T})$ 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 $\Sigma \subseteq \mathsf{sig}(\mathcal{T})$. 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 [Ghilardi *et al.*, 2006]:

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-use* 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 [Wilke, 2001] point the way towards the extension of $\mathcal{ALC}$ with fixpoint operators.

References

[Baader *et al.*, 2003] F. Baader, D. Calvanese, D. McGuiness, D. Nardi, and P. Patel-Schneider. *The Description Logic Handbook*. Cambridge University Press, 2003.

[D'Agostino and Hollenberg, 1998] G. D'Agostino and M. Hollenberg. Uniform interpolation, automata and the modal μ-calculus. In *AiML, Vol. 1*, pages 73–84. CSLI Publications, 1998.

[Eiter *et al.*, 2006] T. Eiter, G. Ianni, R. Schindlauer, H. Tompits, and K. Wang. Forgetting in managing rules and ontologies. In *Web Intelligence*, pages 411–419, 2006.

[French, 2006] T. French. *Bisimulation quantifiers for modal logics*. PhD thesis, University of Western Australia, 2006.

[Ghilardi *et al.*, 2006] S. Ghilardi, C. Lutz, and F. Wolter. Did I damage my ontology? A case for conservative extensions in DL. In *Proc. of KR06*, pages 187–197. AAAI Press, 2006.

[Goranko and Otto, 2007] V. Goranko and M. Otto. Model theory of modal logic. In *Handbook of Modal Logic*. Elsevier, 2007.

[Grau and Motik, 2010] B. Cuenca Grau and B. Motik. Pushing the limits of reasoning over ontologies with hidden content. In *Proc. of KR10*, AAAI Press, 2010.

[Herzig and Mengin, 2008] A. Herzig and J. Mengin. Uniform interpolation by resolution in modal logic. In *Proc. of JELIA08*, vol. 5293 of *LNCS*, pages 219–231. Springer, 2008.

[Konev *et al.*, 2009] 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.

[Kontchakov *et al.*, 2010] R. Kontchakov, F. Wolter, and M. Zakharyaschev. Logic-based ontology comparison and module extraction, with an application to DL-Lite. *Artificial Intelligence*, 174(15):1093–1141, 2010.

[Lang *et al.*, 2003] J. Lang, P. Liberatore, and P. Marquis. Propositional independence: formula-variable independence and forgetting. *J. of Artificial Intelligence Research*, 18:391–443, 2003.

[Reiter and Lin, 1994] R. Reiter and F. Lin. Forget it! In *Proc. of AAAI Fall Symposium on Relevance*, pages 154–159, 1994.

[ten Cate *et al.*, 2006] B. ten Cate, W. Conradie, M. Marx, and Y. Venema. Definitorially complete description logics. In *Proc. of KR06*, pages 79–89. AAAI Press, 2006.

[Visser, 1996] A. Visser. Uniform interpolation and layered bisimulation. In *Proc. of Gödel96*, vol.6 of LNL, pages 139–164. Springer Verlag, 1996.

[Wang *et al.*, 2008] 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*, vol. 5823 of LNCS, pages 666–681. Springer, 2009.

[Wang *et al.*, 2009b] Z. Wang, K. Wang, R.W. Topor, J.Z. Pan, and G. Antoniou. Uniform interpolation for $\mathcal{ALC}$ revisited. In *Proc. of ACAI09*, vol. 5866 of LNCS, pages 528–537. Springer, 2009.

[Wang *et al.*, 2010] 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.

[Wilke, 2001] T. Wilke. Alternating tree automata, parity games, and modal μ-calculus. *Bulletin of the Belg. Math. Soc*, 8(2), 2001.