Computation of Interpolants for Description Logic Concepts in Hard Cases

Jean Christoph Jung[1], Jędrzej Kołodziejski[2] and Frank Wolter[3]

[1]TU Dortmund University, Germany
[2]TU Dortmund University, Germany
[3]University of Liverpool, UK

Abstract

While the computation of Craig interpolants for description logics (DLs) with the Craig Interpolation Property (CIP) is well understood, very little is known about the computation and size of interpolants for DLs without CIP or if one aims at interpolating concepts in a weaker DL than the DL of the input ontology and concepts. In this paper, we provide the first elementary algorithms computing (i) $\mathcal{ALC}$ interpolants between $\mathcal{ALC}$-concepts under $\mathcal{ALCH}$-ontologies and (ii) $\mathcal{ALC}$ interpolants between $\mathcal{ALCQ}$-concepts under $\mathcal{ALCQ}$-ontologies. The algorithms are based on recent decision procedures for interpolant existence. We also observe that, in contrast, uniform depth restricted interpolants might be of non-elementary size.

Keywords

Interpolation, Separation, Description Logic

1. Introduction

Interpolants between description logic (DL) concepts have found many applications. For instance, they can be used as explicit concept definitions or referring expressions, as explanations for concept inclusions, as rewritings of queries, and as separating concepts in the context of concept learning [1, 2, 3, 4, 5]. The computation of interpolants has been investigated extensively, both by the DL community [6, 7, 8, 9] but also in modal logic and related fragments of FO [10, 11, 12]. We quickly remind the reader how this is done: A Craig interpolant between C and D is a concept E in the shared signature of C and D such that $\models C \sqsubseteq E$ and $\models E \sqsubseteq D$ (for simplicity we drop the ontology). A DL has the Craig Interpolation Property (CIP), if the existence of such an interpolant follows from $\models C \sqsubseteq D$. DLs such as $\mathcal{ALC}$, $\mathcal{ALCQ}$, and $\mathcal{ALCI}$ have the CIP [6]. Then, an interpolant E can typically be extracted from a proof of $\models C \sqsubseteq D$ (or, equivalently, of non-satisfiability of $C \sqcap \neg D$) in standard calculi in the field such as tableau, the chase, sequent calculi, or type elimination [6, 7, 8, 11, 13].

The situation is very different for DLs that do not enjoy the CIP or if one is interested in interpolating concepts in a weaker DL than the concepts used in the inclusion. In this case, the existence of an interpolating concept does not follow from the validity of the inclusion and extracting interpolating concepts from proofs becomes much harder. In fact, very little is known about how this could be done and research has so far focused on deciding the existence of interpolants rather than constructing them [14, 15, 16]. It is worth noting, however, that for extensions of $\mathcal{EL}$, the chase can be used to compute interpolants even without CIP [7].

It is well known that Craig interpolants of $\mathcal{ALC}$ concept inclusions under $\mathcal{ALCH}$ ontologies do not necessarily exist [17, 6] and that not every $\mathcal{ALCQ}$ concept inclusion has an $\mathcal{ALC}$ interpolant (take $\models C \sqsubseteq C$ for any $\mathcal{ALCQ}$ concept C not equivalent to an $\mathcal{ALC}$ concept). Existence of $\mathcal{ALC}$ interpolants in these settings is, however, decidable [15, 18]. To explain the proof, assume that $\models C \sqsubseteq D$ and let Σ be any signature (again we drop the ontology for simplicity). It is known that an interpolating $\mathcal{ALC}(\Sigma)$ concept exists if no Σ-bisimilar nodes satisfying C and $\neg D$ exist. Hence it suffices to decide whether a

DL 2025: 38th International Workshop on Description Logics, September 3–6, 2025, Opole, Poland

jean.jung@tu-dortmund.de (J. C. Jung); jedrzej.kolodziejski@tu-dortmund.de (J. Kołodziejski); wolter@liverpool.ac.uk (F. Wolter)
0000-0002-4159-2255 (J. C. Jung); 0000-0001-5008-9224 (J. Kołodziejski); 0000-0002-4470-606X (F. Wolter)

pair of concepts is satisfiable in Σ-bisimilar nodes. It turns out that to decide this problem it is crucial to decide the more general problem whether a *set* of concepts (and not just a pair) is satisfiable in mutually Σ-bisimilar nodes. By completing concepts to types containing them, it suffices to decide the latter problem for sets of types, often called *mosaics*. In fact, the decision algorithms in [15, 18] use mosaics and generalize the well-known *type elimination procedures* deciding satisfiability of concepts to *mosaic elimination procedures* deciding Σ-bisimilar satisfiability.

Mosaic elimination procedures decide the existence of interpolants, but they do not construct any interpolants. The aim of this paper is to give the first elementary algorithms constructing $\mathcal{ALC}$ interpolants whenever they exist under ontologies in $\mathcal{ALCH}$ and under ontologies and concept inclusions in $\mathcal{ALCQ}$. Our algorithms are not restricted to computing Craig interpolants, but work for arbitrary signatures. The idea of the algorithms is to run the mosaic elimination procedures discussed above and construct, in addition and inductively, for each eliminated mosaic entailed $\mathcal{ALC}(\Sigma)$ concepts that witness non Σ-bisimilar satisfiability of its types. The witness concepts we propose are not aggregated at each step, but are polyadic in the sense that we define, for any set T of concepts (types in the case of mosaics) which are not satisfiable in Σ-bisimilar nodes, for each $C \in T$ an $\mathcal{ALC}(\Sigma)$ concept $\mathsf{Sep}(C)$ such that the following holds:

- $\models C \sqsubseteq \mathsf{Sep}(C)$ for all $C \in T$;
- $\models \bigsqcap_{C \in T} \mathsf{Sep}(C) \sqsubseteq \bot$.

The concept $\mathsf{Sep}(C)$ constructed for $T = \{C, \neg D\}$ is then the desired interpolant. We note that an earlier attempt to construct interpolants while running a mosaic elimination procedure without using polyadic separators does not work as stated [15]. Hence one main contribution of this paper is to correct that proof. Our second main contribution is to show that our approach also works in the case of $\mathcal{ALCQ}$ ontologies.

2. Preliminaries

We first introduce the syntax and semantics of the basic description logics $\mathcal{ALC}$, $\mathcal{ALCH}$, and $\mathcal{ALCQ}$ and introduce some model theory. We refer the reader to [19] for a comprehensive introduction to description logics. Let $\mathsf{N_C}$, and $\mathsf{N_R}$ be mutually disjoint and countably infinite sets of *concept*, and *role names*. An $\mathcal{ALCQ}$ *concept* is defined according to the syntax rule

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

where A ranges over concept names, r over role names, and $n \geq 0$. We use the standard abbreviations $\exists r.C$ for $(\geq 1\ r.C)$, $\forall r.C$ for $\neg\exists r.\neg C$, $C \sqcup D$ for $\neg(\neg C \sqcap \neg D)$, and $C \to D$ for $\neg C \sqcup D$. An $\mathcal{ALC}$ *concept* is an $\mathcal{ALCQ}$ concept in which for every subformula $(\geq 1\ r.C)$, n is actually 1. An $\mathcal{ALCQ}$ *concept inclusion ($\mathcal{ALCQ}$ CI)* takes the form $C \sqsubseteq D$ for $\mathcal{ALCQ}$ concepts C and D. $\mathcal{ALC}$ *concept inclusions* are defined accordingly. An $\mathcal{ALCQ}$ *ontology* is a finite set of $\mathcal{ALCQ}$ CIs. An $\mathcal{ALCH}$ *ontology* is a finite set of $\mathcal{ALC}$ concept inclusions and *role inclusions (RIs)* $r \sqsubseteq s$ where r, s are role names from $\mathsf{N_R}$. The *size* of a (finite) syntactic object X, denoted $\|X\|$, is the number of symbols needed to represent it as a word, and the *role depth* of a concept is the maximal nesting depth of concept constructors $(\geq n\ r.C)$.

As usual, the semantics is defined in terms of *interpretations* $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, where $\Delta^{\mathcal{I}}$ is a non-empty set, called *domain* of $\mathcal{I}$, and $\cdot^{\mathcal{I}}$ is a function mapping every $A \in \mathsf{N_C}$ to a subset of $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$ and every $r \in \mathsf{N_R}$ to a subset of $r^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$. The *extension $C^{\mathcal{I}}$ of a concept C in $\mathcal{I}$* is defined as usual. An interpretation $\mathcal{I}$ *satisfies* a CI $C \sqsubseteq D$ if $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$ and an RI $r \sqsubseteq s$ if $r^{\mathcal{I}} \subseteq s^{\mathcal{I}}$. We say that $\mathcal{I}$ is a *model* of an ontology $\mathcal{O}$ if it satisfies all inclusions in it. A concept C is *satisfiable under ontology $\mathcal{O}$* if there is a model $\mathcal{I}$ of $\mathcal{O}$ with $C^{\mathcal{I}} \neq \emptyset$. Moreover, C *is subsumed by another concept D under $\mathcal{O}$* if $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$ in every model $\mathcal{I}$ of $\mathcal{O}$. We write $\mathcal{O} \models C \sqsubseteq D$ in this case.

We next introduce the studied notions and associated problems. A *signature* is a set Σ of concept and role names. An $\mathcal{ALC}(\Sigma)$ *concept* is an $\mathcal{ALC}$ concept that uses only concept and role names from Σ.

Let $\mathcal{L}, \mathcal{L}'$ be DLs, and let us fix an $\mathcal{L}$ ontology $\mathcal{O}$, $\mathcal{L}$ concepts C, D, and a signature Σ. Then, an $\mathcal{L}'(\Sigma)$ *interpolant for* $\mathcal{O} \models C \sqsubseteq D$ is an $\mathcal{L}'(\Sigma)$ concept E with $\mathcal{O} \models C \sqsubseteq E$ and $\mathcal{O} \models E \sqsubseteq D$. The associated decision problem of $\mathcal{L}'(\Sigma)$ *interpolant existence* over $\mathcal{L}$ ontologies and concepts has been recently studied and shown decidable [15, 18]. The decision procedures are based on elegant model-theoretic characterizations of interpolant existence in terms of bisimulations, which we introduce next. A relation $Z \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{J}}$ is a Σ-*bisimulation* between interpretations $\mathcal{I}$ and $\mathcal{J}$ if the following conditions are satisfied for all $(d, e) \in Z$:

Atom for all concept names $A \in \Sigma$: $d \in A^{\mathcal{I}}$ iff $e \in A^{\mathcal{J}}$,

Back for all role names $r \in \Sigma$ and all $(d, d') \in r^{\mathcal{I}}$, there is $(e, e') \in r^{\mathcal{J}}$ such that $(d', e') \in Z$,

Forth for all role names $r \in \Sigma$ and all $(e, e') \in r^{\mathcal{J}}$, there is $(d, d') \in r^{\mathcal{I}}$ such that $(d', e') \in Z$.

A *pointed interpretation* is a pair $\mathcal{I}, d$ with $\mathcal{I}$ an interpretation and $d \in \Delta^{\mathcal{I}}$. We write $\mathcal{I}, d \sim_{\mathcal{ALC},\Sigma} \mathcal{J}, e$ and call $\mathcal{I}, d$ and $\mathcal{J}, e$ Σ-*bisimilar* if there exists an Σ-bisimulation Z such that $(d, e) \in Z$. We say that $\mathcal{ALCQ}$ concepts C_1, C_2 are *jointly* $\sim_{\mathcal{ALC},\Sigma}$-*consistent under* $\mathcal{O}$ if there are models $\mathcal{I}_1, \mathcal{I}_2$ of $\mathcal{O}$ and elements $d_i \in C_i^{\mathcal{I}_i}$ for $i = 1, 2$ with $\mathcal{I}_1, d_1 \sim_{\mathcal{ALC},\Sigma} \mathcal{I}_2, d_2$. We have the following characterization:

Lemma 1. *Let $\mathcal{L} \in \{\mathcal{ALCH}, \mathcal{ALCQ}\}$, $\mathcal{O}$ be an $\mathcal{L}$ ontology, C, D be $\mathcal{L}$-concepts, and Σ be a signature. Then the following are equivalent:*

1. *there is an $\mathcal{ALC}(\Sigma)$ interpolant for $\mathcal{O} \models C \sqsubseteq D$;*
2. *$C, \neg D$ are not jointly $\sim_{\mathcal{ALC},\Sigma}$-consistent under $\mathcal{O}$.*

The proof of Lemma 1 is based on the fact that Σ-bisimulations capture the expressive power of $\mathcal{ALC}(\Sigma)$ concepts, and crucially relies on the use of compactness. In particular, it is not constructive in the sense that in the proof of implication $2 \Rightarrow 1$, no interpolant is constructed. We study here the associated computation problems, that is, compute the interpolants if they exist. A notion dual to the notion of an interpolant is that of a *separator*. Given concepts C, D, E we call E a separator for C, D if $C \sqsubseteq E$ and $E \sqsubseteq \neg D$. Clearly, E is a separator for C, D iff it is an interpolant for $C, \neg D$. Thus, the problems of finding interpolants and separators for a given pair of concepts are algorithmically equivalent. We will switch between these two perspectives depending on which one is more convenient in a given context.

3. Role Inclusions

In this section, we are concerned with computing $\mathcal{ALC}$ interpolants of concept inclusions under $\mathcal{ALCH}$ ontologies. We start with an example that illustrates the failure of the computation algorithm given in [15].

Example 2. *Fix $k \geq 1$, $\mathcal{O} = \{r \sqsubseteq s_i, s_i \sqsubseteq s_i' \mid i \leq k\}$, $\Sigma = \{s_i', A_i \mid i \leq k\}$, and*

$$C = \exists r.B \sqcap \forall r.(B \to \bigsqcup_{i \leq k} A_i) \qquad \text{and} \qquad D = \bigsqcap_{i \leq k} \forall s_i.\neg A_i.$$

We show that C, D are not jointly $\sim_{\mathcal{ALC},\Sigma}$-consistent under $\mathcal{O}$. Indeed, if C, D are jointly $\sim_{\mathcal{ALC},\Sigma}$-consistent, then all concepts in $S = \{B \sqcap (B \to \bigsqcup_{i \leq k} A_i)\} \cup \{\neg A_i \mid i \leq k\}$ are satisfied in mutually Σ-bisimilar nodes, which is clearly not the case. By Lemma 1, there is an $\mathcal{ALC}(\Sigma)$ interpolant for $\mathcal{O} \models C \sqsubseteq \neg D$. For instance, $E = \bigsqcup_{i \leq k} \exists s_i'.A_i$ is an $\mathcal{ALC}(\Sigma)$ interpolant. The algorithm from [15], however, computes a concept of shape $\exists s_i'.E$ for a single $i \leq k$ and one can easily see that a concept of this shape cannot serve as an interpolant. The mistake in the algorithm is a confusion in the quantifier order in the assumptions of the interpolant construction.

We first show how a natural idea for computing interpolants, which works in several other settings, fails to compute elementary sized interpolants in the presence of role inclusions. Next we provide an algorithm which does compute elementary interpolants.

A natural idea to compute interpolants could be to show first that, if there is an $\mathcal{ALC}(\Sigma)$ interpolant for $\mathcal{O} \models C_0 \sqsubseteq D_0$, then there is one of small role depth n, and then use the strongest $\mathcal{ALC}$ consequence of C_0 of this role depth n. Let $n \geq 0$. A (Σ, n)-*uniform interpolant of* C_0 *under* $\mathcal{O}$ is an $\mathcal{ALC}(\Sigma)$ concept U such that $\mathcal{O} \models C_0 \sqsubseteq U$ and $\mathcal{O} \models U \sqsubseteq E$ for every $\mathcal{ALC}(\Sigma)$ concept E of role depth at most n with $\mathcal{O} \models C_0 \sqsubseteq E$. A (Σ, n)-uniform interpolant for C_0 under $\mathcal{O}$ always exists, and can be used as an interpolant for $\mathcal{O} \models C_0 \sqsubseteq D$ whenever an $\mathcal{ALC}(\Sigma)$ interpolant for $\mathcal{O} \models C_0 \sqsubseteq D$ of role depth n exists. This idea has been used to compute elementary sized modal logic interpolants of μ-calculus formulae [20] and it follows from its proof that it applies to computing interpolants under $\mathcal{ALC}$ ontologies. We actually conjecture that it works for the majority of DLs enjoying the CIP, but we leave an elaboration for future work. Unfortunately, contrary to these settings, in our case (Σ, n)-uniform interpolants need not be elementary in n, and consequently do not lead to elementary sized interpolants. In what follows we denote by Tower the iterated exponential function, that is, $\textsc{Tower}(0) = 1$ and $\textsc{Tower}(n+1) = 2^{\textsc{Tower}(n)}$.

Theorem 3. *There is an $\mathcal{ALCH}$ ontology $\mathcal{O}$, an $\mathcal{ALC}$ concept C_0, and signature Σ such that there is no (Σ, n)-uniform interpolant of C_0 under $\mathcal{O}$ smaller than* $\textsc{Tower}(n-2)$.

Proof. Consider the $\mathcal{ALCH}$ ontology $\mathcal{O} = \{r \sqsubseteq s, r \sqsubseteq s'\}$, the concept $C_0 = \exists r.\top$ and $\Sigma = \{s, s'\}$. We claim that no concept U of size smaller than $\textsc{Tower}(n-2)$ is a (Σ, n)-uniform interpolant for C_0 under $\mathcal{O}$. Assume towards contradiction that there is such an U. Observe that for every $C \in \mathcal{ALC}(\Sigma)$ of depth $n-1$, $\forall s.C \to \exists s'.C$ is an $\mathcal{ALC}(\Sigma)$ concept of depth n and $\mathcal{O} \models C_0 \sqsubseteq (\forall s.C \to \exists s'.C)$. Hence $\mathcal{O} \models U \sqsubseteq (\forall s.C \to \exists s'.C)$. Consider all trees of depth $n-1$, choose one for every equivalence class of Σ-bisimulation and denote the set of all these chosen trees by $\mathcal{T}$. We have $\textsc{Tower}(n-1) \leq |\mathcal{T}|$ and $\|U\| < \textsc{Tower}(n-2)$. Thus, by the pigeonhole principle there are two different $\mathcal{I}_1, \mathcal{I}_2 \in \mathcal{T}$ whose respective roots d_1, d_2 satisfy exactly the same sub-concepts of U. Every two trees in $\mathcal{T}$ are distinguished by some $D \in \mathcal{ALC}(\Sigma)$ of depth $n-1$, so let us pick D such that $d_1 \in D^{\mathcal{I}_1}$ but $d_2 \notin D^{\mathcal{I}_2}$. We claim that $\mathcal{O} \not\models U \sqsubseteq \forall s.D \to \exists s'.D$, which contradicts that U is an (Σ, n)-uniform interpolant. This is witnessed by an interpretation $\mathcal{J}$ constructed as follows. First take the disjoint union of $\mathcal{I}_1, \mathcal{I}_2$. Then take two fresh points e_1, e_2, and add edges $e_1 \xrightarrow{r} d_1, e_1 \xrightarrow{s} d_1, e_1 \xrightarrow{s'} d_1$ and $e_2 \xrightarrow{s} d_1, e_2 \xrightarrow{s'} d_2$. Since $\mathcal{O} \models C_0 \sqsubseteq U$ and C_0 is true at e_1 we have $e_1 \in U^{\mathcal{J}}$. This implies $e_2 \in U^{\mathcal{J}}$ because d and d' satisfy the same subconcepts of U. But $\forall s.D \to \exists s'.D$ is false at e_2, a contradiction.

□

On the positive side, we show the following second main result.

Theorem 4. *Let $\mathcal{O}$ be an $\mathcal{ALCH}$ ontology, C_0, D_0 $\mathcal{ALC}$ concepts, and Σ be a signature. Then, if there is an $\mathcal{ALC}(\Sigma)$ interpolant for $\mathcal{O} \models C_0 \sqsubseteq D_0$, we can construct the DAG representation of such an interpolant in time double exponential in $\|\mathcal{O}\| + \|C_0\| + \|D_0\|$.*

The proof is by extending a known mosaic elimination procedure for deciding joint $\sim_{\mathcal{ALC},\Sigma}$-consistency for input ontology and concepts formulated in $\mathcal{ALCH}$ [15]. We present a slight simplification of the original procedure, as we require it only for a restricted setting.

Let us fix an $\mathcal{ALCH}$ ontology $\mathcal{O}$, $\mathcal{ALC}$ concepts C_0, D_0, and a signature Σ. We denote with $\mathsf{sub}(\mathcal{O}, C_0, D_0)$ the set of subconcepts that occur in $\mathcal{O}, C_0, D_0$, closed under single negation. A *type for* $\mathcal{O}$ is any subset of $\mathsf{sub}(\mathcal{O}, C_0, D_0)$ *realizable* in a model of $\mathcal{O}$, that is, any set $t \subseteq \mathsf{sub}(\mathcal{O}, C_0, D_0)$ such that there is a model $\mathcal{I}$ of $\mathcal{O}$ and element $d \in \Delta^{\mathcal{I}}$ with $t = \mathsf{tp}_{\mathcal{I}}(d)$ where:

$$\mathsf{tp}_{\mathcal{I}}(d) = \{E \in \mathsf{sub}(\mathcal{O}, C_0, D_0) \mid d \in E^{\mathcal{I}}\}.$$

We often treat a type t as the conjunction of all concepts it contains, which allows us to write, for instance, $\mathcal{O} \models t \sqsubseteq D$. A *mosaic for* $\mathcal{O}$ is a set T of types for $\mathcal{O}$. We say that a type t is a *completion* of a

concept $C \in \mathsf{sub}(\mathcal{O}, C_0, D_0)$ if $C \in t$, and a mosaic T is a *completion* of a set $\mathcal{C} \subseteq \mathsf{sub}(\mathcal{O}, C_0, D_0)$ of concepts if T contains a completion of every $C \in \mathcal{C}$.

Intuitively, a mosaic T describes a collection of elements in an interpretation $\mathcal{I}$ which realize precisely the types in T and are mutually Σ-bisimilar. Naturally, not every set of types can be realized in this way, and we use a mosaic elimination procedure to determine which can. We write $t \rightsquigarrow_r t'$ if an element of type t' is a viable r-successor of an element of type t, that is, $\{C \mid \forall r.C \in t\} \subseteq t'$. We will denote $\{C \mid \forall r.C \in t\} = t_{/r}$. We write $T \rightsquigarrow_r T'$ if for every $t \in T$, there is $t' \in T'$ with $t \rightsquigarrow_r t'$. Let $\mathcal{M}$ be a set of mosaics. A mosaic $T \in \mathcal{M}$ is *bad* if it violates one of the following conditions:

(Atomic Consistency) for every $t, t' \in T$ and $A \in \Sigma$, $A \in t$ iff $A \in t'$;

(Existential Saturation) for every $t \in T$ and $\exists r.C \in t$, there is $T' \in \mathcal{M}$ such that (a) $C \in t'$ for some $t' \in T'$ with $t \rightsquigarrow_r t'$ and (b) if $\mathcal{O} \models r \sqsubseteq s$ for some $s \in \Sigma$, then $T \rightsquigarrow_s T'$.

Along the lines of the proof of Lemma 6.5 in [15] one can show Lemma 5 below, see the appendix for a proof sketch. The original Lemma 6.5 works with pairs of mosaics which is necessary for DLs that are not preserved under disjoint unions such as $\mathcal{ALCO}$.

Lemma 5. *C_0 and D_0 are jointly $\sim_{\mathcal{ALC},\Sigma}$-consistent under $\mathcal{O}$ iff there is a set $\mathcal{M}^*$ of mosaics that does not contain a bad mosaic and such that there is $T \in \mathcal{M}^*$ and $t_1, t_2 \in T$ with $C_0 \in t_1$ and $D_0 \in t_2$.*

It is a consequence of Lemma 5 that joint $\sim_{\mathcal{ALC},\Sigma}$-consistency under $\mathcal{ALCH}$ ontologies can be decided in double exponential time. Indeed, an $\mathcal{M}^*$ as in Lemma 5 can be found (if it exists) by exhaustively eliminating bad mosaics from the set of all mosaics. Since the set of all mosaics is of double exponential size, and each round of the elimination procedure can be performed in time polynomial in the size of the current set of mosaics, the upper bound follows. By the link to interpolant existence provided in Lemma 1, also $\mathcal{ALC}(\Sigma)$ interpolant existence is decidable in double exponential time.

Our aim is to extend the described mosaic elimination procedure by computing, for each type in an eliminated mosaic, its "contribution" to the elimination. To formalize this we introduce a polyadic notion of a separator reflecting the fact that a mosaic may contain more than two types. Assume a set $\mathcal{C}$ of concepts. An $\mathcal{ALC}(\Sigma)$ *separator for* $\mathcal{C}$ is a function Sep from $\mathcal{C}$ to $\mathcal{ALC}(\Sigma)$-concepts such that:

- $\mathcal{O} \models C \sqsubseteq \mathsf{Sep}(C)$ for every $C \in \mathcal{C}$;
- $\mathcal{O} \models \bigsqcap_{C \in \mathcal{C}} \mathsf{Sep}(C) \sqsubseteq \bot$.

We call $\mathcal{C}$ $\mathcal{ALC}(\Sigma)$*-separable* if there is an $\mathcal{ALC}(\Sigma)$ separator for $\mathcal{C}$. We will use the following lemma which connects separation of concepts with separation of their completions.

Lemma 6. *Assume a set $\mathcal{C} \subseteq \mathsf{sub}(\mathcal{O}, C_0, D_0)$ of concepts. The following are equivalent:*

1. *the set $\mathcal{C}$ is $\mathcal{ALC}(\Sigma)$-separable;*
2. *every completion T of $\mathcal{C}$ is $\mathcal{ALC}(\Sigma)$-separable.*

Proof. The implication (1) $\Longrightarrow$ (2) is straightforward because a separator for $\mathcal{C}$ is a separator for every completion T of $\mathcal{C}$.

For the other implication (2) $\Longrightarrow$ (1) assume that for every completion of $\mathcal{C}$ we have a separator. Let $\mathsf{cp}(\mathcal{C})$ denote the set of all functions from $\mathcal{C}$ to types which map every $C \in \mathcal{C}$ to one of its completions. The image $f[\mathcal{C}]$ of every such function $f \in \mathsf{cp}(\mathcal{C})$ is a completion of $\mathcal{C}$, and thus, by assumption, is separated by some Sep_f.

We define a separator Sep for $\mathcal{C}$ by setting, for every $C \in \mathcal{C}$:

$$\mathsf{Sep}(C) = \bigsqcup_{\substack{t \text{ completion} \\ \text{of } C}} \bigsqcap_{\substack{f \in \mathsf{cp}(\mathcal{C}) \\ f(C) = t}} \mathsf{Sep}_f(C) \qquad (*)$$

To prove $\mathcal{O} \models C \sqsubseteq \mathsf{Sep}(C)$ assume $\mathcal{I} \models \mathcal{O}$ and $d \in C^{\mathcal{I}}$. Let t be the type $\mathsf{tp}_{\mathcal{I}}(d)$ of d. Clearly, t is a completion of C. By assumption, for every $f \in \mathsf{cp}(\mathcal{C})$ and $C \in \mathcal{C}$ we have $\mathcal{O} \models f(C) \sqsubseteq \mathsf{Sep}_f(C)$.

Hence, for every $f \in \mathsf{cp}(\mathcal{C})$ such that $f(C) = t$ we get $\mathcal{O} \models t \sqsubseteq \mathsf{Sep}_f(C)$. It follows that $d \in (\mathsf{Sep}(C))^{\mathcal{I}}$.

It remains to show that $\mathcal{O} \models \bigsqcap_{C \in \mathcal{C}} \mathsf{Sep}(C) \sqsubseteq \bot$. Assume towards contradiction an interpretation $\mathcal{I} \models \mathcal{O}$ with $d \in (\bigsqcap_{C \in \mathcal{C}} \mathsf{Sep}(C))^{\mathcal{I}}$. By definition of Sep, for every C there is a completion t_C such that d satisfies:

$$\bigsqcap_{\substack{f \in \mathsf{cp}(\mathcal{C}) \\ f(C) = t_C}} \mathsf{Sep}_f(C).$$

Consider a function $f \in \mathsf{cp}(\mathcal{C})$ defined as $f(C) = t_C$ for every C. It follows that d satisfies $\mathsf{Sep}_f(C)$ for every C. This contradicts the assumption that $\{\mathsf{Sep}_f(C) \mid C \in \mathcal{C}\}$ is a separator for the image of f and as such is inconsistent. □

We inductively define separators for each eliminated mosaic. Recall that there are two ways a mosaic T can be eliminated: the base case when T violates atomic consistency, and the inductive case when T violates existential saturation. We look at these cases in turn.

Inductive Base. If T violates atomic consistency then there is a concept name $A \in \Sigma$ and types $t, t' \in T$ with $A \in t$ and $\neg A \in t'$. Let $\mathsf{Sep}(t)$ be A if $A \in t$ and $\neg A$ otherwise. It follows that $\mathcal{O} \models t \sqsubseteq \mathsf{Sep}(t)$ for all t, and $\bigsqcap_{t \in T} \mathsf{Sep}(t) \sqsubseteq \bot$.

Inductive Step. Denote the current set of mosaics by $\mathcal{M}$ and assume a mosaic $T \in \mathcal{M}$ is eliminated because it violates existential saturation. This means that there are $t \in T$ and $\exists r.C \in t$ such that whenever $T' \in \mathcal{M}$ satisfies (i) $T \rightsquigarrow_s T'$ for all $\mathcal{O} \models r \sqsubseteq s$, and (ii) contains some $t' \in T'$ with $t \rightsquigarrow_r t'$ and $C \in t'$ then $T' \notin \mathcal{M}$. Consider the set:

$$\mathcal{D} = \{t'_{/s} \mid t' \in T, s \in \Sigma, \text{ and } \mathcal{O} \models r \sqsubseteq s\} \cup \{\{C\} \cup t_{/r}\}.$$

It follows that every completion T' of $\mathcal{D}$ was already eliminated from $\mathcal{M}$: left and right part of the union correspond to parts (i) and (ii) of the violated condition. Lemma 6 provides us with a separator $\mathsf{Sep}_{\mathcal{D}}$ for $\mathcal{D}$. We use $\mathsf{Sep}_{\mathcal{D}}$ to get a separator Sep for T as follows. We put:

$$\mathsf{Sep}(t') = \bigsqcap_{\substack{\mathcal{O} \models r \sqsubseteq s, \\ s \in \Sigma}} \forall s.\mathsf{Sep}_{\mathcal{D}}(t'_{/s}) \quad \text{and} \quad \mathsf{Sep}(t) = \neg \bigsqcap_{t' \neq t} \mathsf{Sep}(t') = \bigsqcup_{t' \neq t} \bigsqcup_{\substack{\mathcal{O} \models r \sqsubseteq s, \\ s \in \Sigma}} \exists s.\neg \mathsf{Sep}_{\mathcal{D}}(t'_{/s}).$$

for every $t' \neq t$.

We claim that Sep separates T. For every $t' \neq t$ we have $\mathcal{O} \models t' \sqsubseteq \mathsf{Sep}(t')$. This follows because for every $\mathcal{O} \models r \sqsubseteq s$ with $s \in \Sigma$ we have $\models t' \sqsubseteq \forall s.t'_{/s}$ and $\mathcal{O} \models t'_{/s} \sqsubseteq \mathsf{Sep}(t'_{/s})$. To show $\mathcal{O} \models t \sqsubseteq \mathsf{Sep}(t)$ assume an interpretation $\mathcal{I} \models \mathcal{O}$ with $d \in t^{\mathcal{I}}$. The point d has an r-child e satisfying $\{C\} \cup t_{/r}$ and hence also $\mathsf{Sep}_{\mathcal{D}}(\{C\} \cup t_{/r})$. By definition of a separator, the image $\mathsf{Sep}_{\mathcal{D}}[\mathcal{D}]$ of $\mathsf{Sep}_{\mathcal{D}}$ is inconsistent. Thus, the fact that d satisfies $\mathsf{Sep}_{\mathcal{D}}(\{C\} \cup t_{/r})$ implies that some concept $E \in \mathsf{Sep}_{\mathcal{D}}[\mathcal{D}]$ other than $\mathsf{Sep}_{\mathcal{D}}(\{C\} \cup t_{/r})$ must be false at d. We therefore have $d \in (\neg \mathsf{Sep}_{\mathcal{D}}(t'_{/s}))^{\mathcal{I}}$ for some $\mathcal{O} \models r \sqsubseteq s$ with $s \in \Sigma$ and some t'. Since for every $s \in \Sigma$ with $\mathcal{O} \models r \sqsubseteq s$ we have $d \in (t_{/s})^{\mathcal{I}}$ and thus $d \in (\mathsf{Sep}_{\mathcal{D}}(t_{/s}))^{\mathcal{I}}$, it follows that $t' \neq t$. This proves that $d \in (\mathsf{Sep}(t))^{\mathcal{I}}$. Note that $\mathsf{Sep}[T]$ is inconsistent by definition: the concept $\mathsf{Sep}(t)$ is just a negated conjunction $\bigsqcap_{t' \neq t} \mathsf{Sep}(t')$ of the rest. This completes the proof that Sep separates T.

This finishes the construction of separators for every eliminated mosaic. To construct the actual interpolant, note that Lemmas 1 and 5 imply that, if there is an $\mathcal{ALC}(\Sigma)$ interpolant for $\mathcal{O} \models C_0 \sqsubseteq D_0$, then all completions of $\{C_0, \neg D_0\}$ have been eliminated. Lemma 6 provides us with an $\mathcal{ALC}(\Sigma)$ separator Sep for $\{C_0, \neg D_0\}$ and it is easy to see that $\mathsf{Sep}(C_0)$ is the sought $\mathcal{ALC}(\Sigma)$ interpolant.

It remains to analyze the DAG size of the constructed separators, which we do by counting the number of sub-formulae used in the constructed separators. On a high-level, we construct one formula for every type in every eliminated mosaic. This formula is of negligible size 1 in the inductive base, so let us analyze the inductive step. This step relies on Lemma 6, and one can see that the construction

in Equation ($*$) uses double exponentially many sub-formulae. It remains to note that the Lemma is invoked only double exponentially often and that the construction of the separator formulae for the just eliminated concept introduces only double exponentially many sub-formulae. This completes the proof of Theorem 4.

We finish the section with some remarks regarding the size of the constructed interpolants. First, we strongly conjecture that there are examples in which the interpolant is forced to have double exponential role depth, so the upper size bound in Theorem 4 is optimal. Second, it is known that the size of DAG representation of interpolants in standard DLs enjoying the CIP is at most exponential [6, Theorem 3.26] and thus there is an exponential gap.

4. Qualified Number Restrictions

We are concerned with computing $\mathcal{ALC}$ interpolants of concept inclusions in $\mathcal{ALCQ}$ under $\mathcal{ALCQ}$ ontologies. We use the same notation for $\mathcal{ALCQ}$ as in the previous section for $\mathcal{ALCH}$, defined in the obvious way. Our first result is that (Σ, n)-uniform $\mathcal{ALC}$ interpolants can be of non-elementary size.

Theorem 7. *There is an $\mathcal{ALCQ}$ concept C_0 and signature Σ such that there is no (Σ, n)-uniform $\mathcal{ALC}$ interpolant of C_0 smaller than* $\textsc{Tower}(n-2)$.

Proof. Take the concept $C_0 = (\leq 1\ r.\top)$, signature $\Sigma = \{r, s, s'\}$, and let $n > 0$. Using $\mathcal{ALC}(\Sigma)$ concepts $\exists r.C \to \forall r.C$ one can show the lower bound in the same way as in the proof of Theorem 3. □

The main result of this section is as follows.

Theorem 8. *Let $\mathcal{O}$ be an $\mathcal{ALCQ}$ ontology, C_0, D_0 $\mathcal{ALC}$ concepts, and Σ be a signature. If there is an $\mathcal{ALC}(\Sigma)$ interpolant for $\mathcal{O} \models C_0 \sqsubseteq D_0$, then there is one of 3-exponential size which can be constructed in 4-exponential time in $\|\mathcal{O}\| + \|C_0\| + \|D_0\|$.*

Fix an $\mathcal{ALCQ}$ ontology $\mathcal{O}$, $\mathcal{ALCQ}$ concepts C_0, D_0, and a signature Σ. Our algorithm for computing interpolants again relies on a mosaic elimination procedure that determines the mosaics for which there is a model $\mathcal{I}$ of $\mathcal{O}$ which realizes the types in T in mutually Σ-bisimilar nodes. To formalize the elimination condition, we need some new notation. Let $m^\bullet \in \mathbb{N}$ be maximal such that $(\geq m^\bullet\ r.C)$ occurs in $\mathsf{sub}(\mathcal{O}, C_0, D_0)$ for some r, C. Let $N^\bullet = \{0, \dots, m^\bullet\} \cup \{\infty\}$, and define $<$ and $+$ on $N^\bullet$ as usual by setting, for instance, $m^\bullet < \infty$ and $k + \infty = \infty$. For a role name r and type t, a *witnessing function* $w_{r,t}$ assigns to every type t' a $w_{r,t}(t') \in N^\bullet$ such that for each $(\geq n\ r.C) \in \mathsf{sub}(\mathcal{O}, C_0, D_0)$, $(\geq n\ r.C) \in t$ iff $\sum_{C \in t'} w_{r,t}(t') \geq n$. If t is realizable, then there exists a witnessing function $w_{r,t}$ for each role name r: take a model $\mathcal{I}$ of $\mathcal{O}$ realizing t in a node d and define

$$w_{r,t}(t') = \begin{cases} n & \text{if } n = |\{d' \in \Delta^{\mathcal{I}} \mid (d, d') \in r^{\mathcal{I}}, t' = \mathsf{tp}_{\mathcal{I}}(d')\}| \leq m^\bullet \\ \infty & \text{otherwise.} \end{cases} \tag{1}$$

Let T be a mosaic, r a role name, and $(w_{r,t})_{t \in T}$ be witnessing functions. To satisfy the types in a mosaic in mutually Σ-bisimilar nodes one must be able to partition, for $r \in \Sigma$, their r-successors into mosaics so that the back- and-forth conditions of Σ-bisimulations hold. Our formalization of this idea follows [18], but we modify the notation for our purposes. Say that a set $\mathcal{S}$ of mosaics is a *mosaic partition for* $(w_{r,t})_{t \in T}$ if one can assign to each t, t' with $t \in T$ and $w_{r,t}(t') > 0$ a non-empty set $a_r(t, t') \subseteq \mathcal{S}$ (intuitively, the mosaics in $\mathcal{S}$ containing t' as an r-successor of t) with $t' \in T'$ for all $T' \in a_r(t, t')$ in such a way that

- for every $T' \in \mathcal{S}$ and $t \in T$, there exists a $t' \in T'$ with $T' \in a_r(t, t')$;
- for all types t, t', $|a_r(t, t')| \leq w_{r,t}(t')$.

Let $\mathcal{M}$ be a set of mosaics. A mosaic $T \in \mathcal{M}$ is *bad* if it violates one of the following conditions:

(Atomic Consistency) for every $t, t' \in T$ and $A \in \Sigma$, $A \in t$ iff $A \in t'$;

(Existential Saturation) for every role name $r \in \Sigma$ there are witnessing functions $(w_{r,t})_{t\in T}$ and a mosaic partition $\mathcal{S} \subseteq \mathcal{M}$ for $(w_{r,t})_{t\in T}$.

The following result is shown in [18] (using slightly different wording):

Lemma 9. *(i) If the condition (Existential Saturation) is satisfied for some $T \in \mathcal{M}$, then this is witnessed by a mosaic partition $\mathcal{S} \subseteq \mathcal{M}$ with $|\mathcal{S}| \leq m^\bullet \times 2^{2|\mathsf{sub}(\mathcal{O},C_0,D_0)|}$.*

(ii) C_0, D_0 are jointly $\sim_{\mathcal{ALC},\Sigma}$-consistent under $\mathcal{O}$ iff there is a set $\mathcal{M}^*$ of mosaics that does not contain a bad mosaic and such that there is $T \in \mathcal{M}^*$ and $t_1, t_2 \in T$ with $C_0 \in t_1$ and $D_0 \in t_2$.

It is a consequence of Lemma 9 that joint $\sim_{\mathcal{ALC},\Sigma}$-consistency under $\mathcal{ALCQ}$ ontologies can be decided in double exponential time. Indeed, an $\mathcal{M}^*$ as in Lemma 9 can be found (if it exists) by exhaustively eliminating bad mosaics from the set of all mosaics. Since the set of all mosaics is of double exponential size, and each round of the elimination procedure can be performed in double exponential time, the upper bound follows. By the link to interpolant existence provided in Lemma 1, also $\mathcal{ALC}(\Sigma)$ interpolant existence is decidable in double exponential time.

We next exploit the elimination procedure to construct interpolants. Similar to the previous Section 3 we compute an $\mathcal{ALC}(\Sigma)$ separator for every eliminated mosaic. It will be convenient to actually compute something slightly stronger. Let $\mathcal{M}$ be a set of mosaics. A function Sep that maps every t in some $T \in \mathcal{M}$ to an $\mathcal{ALC}(\Sigma)$ concept $\mathsf{Sep}(t)$ is called *general $\mathcal{ALC}(\Sigma)$ separator for $\mathcal{M}$* if for every $T \in \mathcal{M}$ the restriction of Sep to T is an $\mathcal{ALC}(\Sigma)$ separator for T.

We compute, by induction, a general $\mathcal{ALC}(\Sigma)$ separator for the set of eliminated mosaics.

Inductive Base. Assume T has been eliminated because atomic consistency is violated. Then there exists $A \in \Sigma$ such that the function Sep_T defined by setting $\mathsf{Sep}_T(t) = A$ if $A \in T$ and $\mathsf{Sep}_T(t) = \neg A$ otherwise, is an $\mathcal{ALC}(\Sigma)$ separator for T. Let $\mathcal{E}_0$ be the set of all mosaics that violate atomic consistency and let Sep_T be defined as above for $T \in \mathcal{E}_0$. Then we obtain a general separator Sep_0 for $\mathcal{E}_0$ by setting $\mathsf{Sep}_0(t) = \bigsqcap_{T\in\mathcal{E}_0} \mathsf{Sep}_T(t)$, for all $t \in T \in \mathcal{E}_0$.

Inductive Step. Assume $\mathcal{E}_n$ is the set of eliminated mosaics and Sep_n is a general separator for $\mathcal{E}_n$. Let $\mathcal{M}_n$ be the set of mosaics that have not yet been eliminated. Setting $\mathsf{Sep}_n(t) = \top$ for types t that do not occur in any mosaic in $\mathcal{E}_n$, we may assume that Sep_n is defined for all types. Let T be the mosaic eliminated in the next step. Then existential saturation is violated in $\mathcal{M}_n$. Note that this implies that one can pick a role name $r \in \Sigma$ such that for every witnessing functions $w_{r,t}$, $t \in T$, and mosaic partition $\mathcal{S}$ for $(w_{r,t})_{t\in T}$ there is an eliminated mosaic $T' \in \mathcal{S} \cap \mathcal{E}_n$. We fix such an r.

Let $\mathcal{T}$ be the set of all types. Denote by $\mathcal{V}^+$ the set of all conjunctions $C = \bigsqcap_{t\in\mathcal{T}} L_t$ with $L_t \in \{\mathsf{Sep}_n(t), \neg\mathsf{Sep}_n(t)\}$. For any nonempty subset $\mathcal{B} \subseteq \mathcal{V}^+$ we set as usual

$$\nabla_r(\mathcal{B}) = (\bigsqcap_{C\in\mathcal{B}} \exists r.C) \sqcap \forall r.(\bigsqcup_{C\in\mathcal{B}} C).$$

Let $\delta_r(t)$ be the disjunction of all $\nabla_r(\mathcal{B})$ such that $t \sqcap \nabla_r(\mathcal{B})$ is satisfiable under $\mathcal{O}$. Observe that $\mathcal{O} \models t \sqsubseteq \delta_r(t)$. Take any $t_0 \in T$ and set $\mathsf{Sep}(t_0) = \neg \bigsqcap_{t\in T\setminus\{t_0\}} \delta_r(t)$ and $\mathsf{Sep}(t) = \delta_r(t)$ for all $t \in T \setminus \{t_0\}$.

Lemma 10. Sep *is an $\mathcal{ALC}(\Sigma)$ separator for T.*

Proof. It suffices to show that $\mathcal{O} \models t_0 \sqsubseteq \mathsf{Sep}(t_0)$, the remaining conditions are trivial. Assume this is not the case. Then we have a model $\mathcal{I}_{t_0}$ of $\mathcal{O}$ and some $d_{t_0} \in \Delta^{\mathcal{I}_{t_0}}$ such that $d_{t_0} \in (t_0 \sqcap (\bigsqcap_{t\in T\setminus\{t_0\}} \delta_r(t))^{\mathcal{I}_{t_0}}$. Pick the (unique) $\mathcal{B} \subseteq \mathcal{V}^+$ such that $d_{t_0} \in \nabla_r(\mathcal{B})^{\mathcal{I}_{t_0}}$. Then $\nabla_r(\mathcal{B})$ is a disjunct of each $\delta_r(t)$ with $t \in T$. Take interpretations $\mathcal{I}_t$ and d_t with $d_t \in (t \sqcap \nabla_r(\mathcal{B}))^{\mathcal{I}_t}$ for $t \in T \setminus \{t_0\}$. We may assume that all $\mathcal{I}_t$, $t \in T$, coincide (otherwise take their disjoint union) and denote it by $\mathcal{I}$. Next define, for $t \in T$, $w_{r,t}(t')$ as in (1). Then $(w_{r,t})_{t\in T}$ is a family of witnessing functions. Let for each $C \in \mathcal{B}$:

$$T_C = \{\mathsf{tp}_{\mathcal{I}}(d) \mid \text{ there is } t \in T \text{ with } (d_t, d) \in r^{\mathcal{I}} \text{ and } d \in C^{\mathcal{I}}\}$$

By the definition of $\mathcal{B}$, $T_C \neq \emptyset$ for every $C \in \mathcal{B}$. Observe that none of the mosaics T_C is in $\mathcal{E}_n$ because otherwise $\mathcal{O} \models \bigsqcap_{t \in T_C} \mathsf{Sep}_n(t) \sqsubseteq \bot$ which is not the case since $\mathsf{Sep}_n(t)$ is a conjunct of C for all $t \in T_C$.

We show that $\{T_C \mid C \in \mathcal{B}\}$ form a mosaic partition for $(w_{r,t})_{t \in T}$, and so derive a contradiction. To this end define $a_r(t, t') \subseteq \{T_C \mid C \in \mathcal{B}\}$ as

$$\{T_C \mid \text{ there is } d \in C^{\mathcal{I}} \text{ with } (d_t, d) \in r^{\mathcal{I}} \text{ and } \mathtt{tp}_{\mathcal{I}}(d) = t'\}$$

To see that $a_r(t, t')$ is as required, first observe that $t' \in T_C$ for any $T_C \in a_r(t, t')$. Next assume that a T_C with $C \in \mathcal{B}$ and $t \in T$ are given. From $d_t \in (t \sqcap \nabla_r(\mathcal{B}))^{\mathcal{I}}$ we obtain a d with $(d_t, d) \in r^{\mathcal{I}}$ and $d \in C^{\mathcal{I}}$. Let $t' = \mathtt{tp}_{\mathcal{I}}(d)$. Then $T_C \in a_r(t, t')$. The condition $|a_r(t, t')| \leq w_{r.t}(t')$ follows directly from the definitions. □

Using Lemma 10 we obtain a general $\mathcal{ALC}(\Sigma)$ separator Sep_{n+1} for $\mathcal{E}_n \cup \{T\}$ by setting $\mathsf{Sep}_{n+1}(t) = \mathsf{Sep}(t) \sqcap \mathsf{Sep}_n(t)$ for all $t \in T$ and $\mathsf{Sep}_{n+1}(t) = \mathsf{Sep}_n(t)$ for all remaining types.

This finishes the construction of separators for every eliminated mosaic. One can now construct the actual interpolants in exactly the same way as in the previous section for $\mathcal{ALCH}$ via Lemma 6. To compute the size of the interpolants, observe that in the construction above $\|\mathsf{Sep}_{n+1}(t)\| \leq \|\mathsf{Sep}_n(t)\| \times 2^{2^{2^{f(m)}}}$ with f a polynomial function and $m = \|\mathcal{O}\| + \|C_0\| + \|D_0\|$. As Sep_n stabilizes after at most double exponentially many elimination steps, $\|\mathsf{Sep}_n\|$ is bound by a 3-exponential function in $\|\mathcal{O}\| + \|C_0\| + \|D_0\|$. This bound remains 3-exponential under DAG representation. The construction of $\mathsf{Sep}_n(t)$ involves satisfiability checks for concepts of 3-exponential size, so overall the interpolant can be constructed in 4-exponential time.

5. Conclusion

We have presented first non-trivial algorithms for computing $\mathcal{ALC}$ interpolants under $\mathcal{ALCH}$ and $\mathcal{ALCQ}$ ontologies, relying on the new notion of polyadic separators tailored to store witnesses for the fact that a mosaic (or set of types) cannot be realized in mutually bisimilar models. Theorems 4 and 8 demonstrate the inherent difficulty of the problem and explain why previously known techniques do not easily apply in the cases of $\mathcal{ALCH}$ and $\mathcal{ALCQ}$. It is worth to note that Theorem 4 can be easily modified to obtain non-elementary lower bounds for the size of uniform interpolants at the concept level in the presence of $\mathcal{ALCH}$ ontologies. These lower bounds, in turn, translate to non-elementary lower bounds for the size of uniform interpolants at the ontology level in $\mathcal{ALCH}$. This implies that the resolution based calculus for computing uniform interpolants of $\mathcal{ALCH}$ ontologies from [21] cannot run in elementary time, answering a question posed by the authors.

In the future, we would like to extend our algorithms to other standard DL constructors. While we believe that this will be rather easy for some constructors, like inverse roles or the universal role, we expect it to be much more involved in other cases such as nominals. In fact, already unifying the two presented algorithms into one for computing $\mathcal{ALC}$ interpolants under $\mathcal{ALCHQ}$ ontologies appears to be challenging. It would be also interesting to analyze our procedures in the ontology-free cases (or with an ontology containing only role inclusions), for which we expect smaller interpolants. From a practical perspective, it would be interesting to extend implemented tableaux algorithms to be able to compute interpolants. Beyond description logics, it would be very interesting to compute interpolants in the guarded and/or the two-variable fragments of first-order logic [14] or in first-order modal logics [16].

Declaration on Generative AI

The authors have not employed any Generative AI tools.

References

[1] B. ten Cate, W. Conradie, M. Marx, Y. Venema, Definitorially complete description logics, in: Proceedings of the 10th International Conference on Principles of Knowledge Representation and Reasoning, KR 2006, AAAI Press, 2006, pp. 79–89. URL: http://www.aaai.org/Library/KR/2006/kr06-011.php.

[2] A. Artale, A. Mazzullo, A. Ozaki, F. Wolter, On free description logics with definite descriptions, in: Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning, KR 2021, 2021, pp. 63–73. URL: https://doi.org/10.24963/kr.2021/7. doi:10.24963/kr.2021/7.

[3] D. Toman, G. E. Weddell, FO Rewritability for OMQ using Beth Definability and Interpolation, in: Proceedings of the 34th International Workshop on Description Logics, DL 2021, CEUR-WS.org, 2021. URL: http://ceur-ws.org/Vol-2954/paper-29.pdf.

[4] J. C. Jung, C. Lutz, H. Pulcini, F. Wolter, Logical separability of labeled data examples under ontologies, Artif. Intell. 313 (2022) 103785. URL: https://doi.org/10.1016/j.artint.2022.103785. doi:10.1016/j.artint.2022.103785.

[5] J. C. Jung, C. Lutz, H. Pulcini, F. Wolter, Separating data examples by description logic concepts with restricted signatures, in: M. Bienvenu, G. Lakemeyer, E. Erdem (Eds.), Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning, KR 2021, Online event, November 3-12, 2021, 2021, pp. 390–399. URL: https://doi.org/10.24963/kr.2021/37. doi:10.24963/KR.2021/37.

[6] B. ten Cate, E. Franconi, I. Seylan, Beth definability in expressive description logics, J. Artif. Intell. Res. 48 (2013) 347–414.

[7] M. Fortin, B. Konev, F. Wolter, Interpolants and Explicit Definitions in Extensions of the Description Logic EL, in: Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning, 2022, pp. 152–162. URL: https://doi.org/10.24963/kr.2022/16. doi:10.24963/kr.2022/16.

[8] T. S. Lyon, J. Karge, Constructive interpolation and concept-based beth definability for description logics via sequents, in: Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24, International Joint Conferences on Artificial Intelligence Organization, 2024, pp. 3484–3492. URL: https://doi.org/10.24963/ijcai.2024/386. doi:10.24963/ijcai.2024/386, main Track.

[9] P. Koopmann, R. A. Schmidt, Uniform interpolation and forgetting for ALC ontologies with aboxes, in: B. Bonet, S. Koenig (Eds.), Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA, AAAI Press, 2015, pp. 175–181. doi:10.1609/AAAI.V29I1.9206.

[10] V. Sofronie-Stokkermans, Interpolation in local theory extensions, Log. Methods Comput. Sci. 4 (2008). URL: https://doi.org/10.2168/LMCS-4(4:1)2008. doi:10.2168/LMCS-4(4:1)2008.

[11] M. Benedikt, B. ten Cate, M. Vanden Boom, Effective interpolation and preservation in guarded logics, ACM Trans. Comput. Log. 17 (2016) 8:1–8:46.

[12] M. Fitting, R. Kuznets, Modal interpolation via nested sequents, Ann. Pure Appl. Log. 166 (2015) 274–305. URL: https://doi.org/10.1016/j.apal.2014.11.002. doi:10.1016/J.APAL.2014.11.002.

[13] B. ten Cate, Lyndon Interpolation for Modal Logic via Type Elimination Sequences, Technical Report, ILLC, Amsterdam, 2022.

[14] J. C. Jung, F. Wolter, Living without beth and craig: Definitions and interpolants in the guarded and two-variable fragments, in: Proceedings of the 36th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2021, IEEE, 2021, pp. 1–14. URL: https://doi.org/10.1109/LICS52264.2021.9470585. doi:10.1109/LICS52264.2021.9470585.

[15] A. Artale, J. C. Jung, A. Mazzullo, A. Ozaki, F. Wolter, Living without beth and craig: Definitions and interpolants in description and modal logics with nominals and role inclusions, ACM Trans. Comput. Log. 24 (2023) 34:1–34:51. URL: https://doi.org/10.1145/3597301. doi:10.1145/3597301.

[16] A. Kurucz, F. Wolter, M. Zakharyaschev, Definitions and (uniform) interpolants in first-order

modal logic, in: Proceedings of the 20th International Conference on Principles of Knowledge Representation and Reasoning, KR 2023, Rhodes, Greece, September 2-8, 2023, 2023, pp. 417–428. URL: https://doi.org/10.24963/kr.2023/41. doi:10.24963/KR.2023/41.

[17] B. Konev, C. Lutz, D. Walther, F. Wolter, Formal properties of modularisation, in: H. Stuckenschmidt, C. Parent, S. Spaccapietra (Eds.), Modular Ontologies: Concepts, Theories and Techniques for Knowledge Modularization, volume 5445 of *Lecture Notes in Computer Science*, Springer, 2009, pp. 25–66. URL: https://doi.org/10.1007/978-3-642-01907-4_3. doi:10.1007/978-3-642-01907-4_3.

[18] L. Kuijer, T. Tan, F. Wolter, M. Zakharyaschev, Separation and definability in fragments of two-variable first-order logic with counting, 2025. URL: https://arxiv.org/abs/2504.20491. arXiv:2504.20491.

[19] F. Baader, I. Horrocks, C. Lutz, U. Sattler, An Introduction to Description Logic, Cambridge University Press, 2017.

[20] J. C. Jung, J. Kolodziejski, Modal separation of fixpoint formulae, in: O. Beyersdorff, M. Pilipczuk, E. Pimentel, K. T. Nguyen (Eds.), 42nd International Symposium on Theoretical Aspects of Computer Science, STACS 2025, March 4-7, 2025, Jena, Germany, volume 327 of *LIPIcs*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025, pp. 55:1–55:20. doi:10.4230/LIPICS.STACS.2025.55.

[21] P. Koopmann, R. Schmidt, Count and forget: Uniform interpolation of shq-ontologies, in: International Joint Conference on Automated Reasoning, 2014. URL: https://api.semanticscholar.org/CorpusID:38753229.

A. Missing Proof Details

Lemma 11. *C_0 and D_0 are jointly $\sim_{\mathcal{ALC},\Sigma}$-consistent under $\mathcal{O}$ iff there is a set $\mathcal{M}^*$ of mosaics that does not contain a bad mosaic and such that there is $T \in \mathcal{M}^*$ and $t_1, t_2 \in T$ with $C_0 \in t_1$ and $D_0 \in t_2$.*

Proof (Sketch). For implication "⇒", suppose that C_0, D_0 are jointly $\sim_{\mathcal{ALC},\Sigma}$-consistent under $\mathcal{O}$, that is, there are models $\mathcal{I}_1, \mathcal{I}_2$ of $\mathcal{O}$ and elements $d_1 \in C_0^{\mathcal{I}_1}$ and $d_2 \in D_0^{\mathcal{I}_2}$ such that $\mathcal{I}_1, d_1 \sim_{\mathcal{ALC},\Sigma} \mathcal{I}_2, d_2$. Since $\mathcal{ALCH}$ is preserved under taking disjoint unions, we can assume without loss of generality that $\mathcal{I}_1 = \mathcal{I}_2 = \mathcal{I}$. We read off a set $\mathcal{M}^*$ of mosaics by taking

$$\mathcal{M}^* = \{\ \{\mathsf{tp}_{\mathcal{I}}(e) \mid \mathcal{I}, d \sim_{\mathcal{ALC},\Sigma} \mathcal{I}, e\}\ \mid d \in \Delta^{\mathcal{I}}\}.$$

It is routine to verify that $\mathcal{M}^*$ satisfies the conditions formulated in Lemma 5.

For implication "⇐", let $\mathcal{M}^*$ be a set of mosaics that does not contain a bad mosaic and such that there is $T^* \in \mathcal{M}^*$ and $t_1, t_2 \in T^*$ with $C \in t_1$ and $D \in t_2$. We construct an interpretation $\mathcal{I}$ as follows:

$$\begin{aligned}
\Delta^{\mathcal{I}} &= \{(t, T) \mid T \in \mathcal{M}^* \text{ and } t \in T\} \\
A^{\mathcal{I}} &= \{(t, T) \in \Delta^{\mathcal{I}} \mid A \in t\} \\
r^{\mathcal{I}} &= \{((t, T), (t', T')) \in \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}} \mid t \rightsquigarrow_r t' \text{ and for all } s \in \Sigma\text{: } ((\mathcal{O} \models r \sqsubseteq s) \Rightarrow T \rightsquigarrow_s T')\}
\end{aligned}$$

One can verify by structural induction that $C \in (t, T)$ iff $(t, T) \in C^{\mathcal{I}}$, for all $C \in \mathsf{sub}(\mathcal{O}, C_0, D_0)$ and $(t, T) \in \Delta^{\mathcal{I}}$. Consequently, $(t_1, T^*) \in C_0^{\mathcal{I}}$ and $(t_2, T^*) \in D_0^{\mathcal{I}}$. Moreover, following relation Z:

$$Z = \{(t, T), (t', T) \mid T \in \mathcal{M}^*\}$$

is a Σ-bisimulation. Since $((t_1, T^*), (t_2, T^*)) \in Z$, we conclude that C_0, D_0 are jointly $\sim_{\mathcal{ALC},\Sigma}$-consistent under $\mathcal{O}$. □