

Fitting Description Logic Ontologies to ABox and Query Examples

Maurice Funk, Marvin Grosser, Carsten Lutz

Leipzig University

ScaDS.AI Center Dresden/Leipzig

{mfunk, grosser, clu}@informatik.uni-leipzig.com

Abstract

We study a fitting problem inspired by ontology-mediated querying: given a collection of positive and negative examples of the form (\mathcal{A}, q) with \mathcal{A} an ABox and q a Boolean query, we seek an ontology \mathcal{O} that satisfies $\mathcal{A} \cup \mathcal{O} \models q$ for all positive examples and $\mathcal{A} \cup \mathcal{O} \not\models q$ for all negative examples. We consider the description logics \mathcal{ALC} and \mathcal{ALCT} as ontology languages and a range of query languages that includes atomic queries (AQs), conjunctive queries (CQs), and unions thereof (UCQs). For all of the resulting fitting problems, we provide effective characterizations and determine the computational complexity of deciding whether a fitting ontology exists. This problem turns out to be CONP -complete for AQs and full CQs and 2EXPTIME -complete for CQs and UCQs. These results hold for both \mathcal{ALC} and \mathcal{ALCT} .

1 Introduction

In many areas of computer science and AI, a fundamental problem is to fit a formal object to a given collection of examples. In inductive program synthesis, for instance, one wants to find a program that complies with a given collection of examples of input-output behavior (Jacindha, Abishek, and Vasuki 2022). In machine learning, fitting a model to a given set of examples is closely linked to PAC-style generalization guarantees (Shalev-Shwartz and Ben-David 2014). And in database research, the traditional query-by-example paradigm asks to find a query that fits a given set of data examples (Li, Chan, and Maier 2015).

In this article, we study the problem of fitting an ontology formulated in a description logic (DL) to a given collection of positive and negative examples. Our concrete setting is motivated by the paradigm of ontology-mediated querying where data is enriched by an ontology that provides domain knowledge, aiming to return more complete answers and to bridge heterogeneous representations in the data (Bienvenu and Ortiz 2015; Xiao et al. 2018). Guided by this application, we use examples that take the form (\mathcal{A}, q) where \mathcal{A} is an ABox (in other words: a database) and q is a Boolean query. We then seek an ontology \mathcal{O} that satisfies $\mathcal{A} \cup \mathcal{O} \models q$ for all positive examples and $\mathcal{A} \cup \mathcal{O} \not\models q$ for all negative examples. It is not a restriction that q is required to be Boolean since our queries may contain individuals from the ABox.

A main application of this ontology fitting problem is to assist with ontology construction and engineering. This is in

the spirit of several other proposals that have the same aim, such as ontology construction and completion using formal concept analysis (Baader et al. 2007; Baader and Distel 2009; Kriegel 2024) and Angluin’s framework of exact learning (Konev et al. 2017), see also the survey (Ozaki 2020). We remark that there is a large literature on fitting DL concepts (rather than ontologies) to a collection of examples, sometimes referred to as concept learning, see for instance (Lehmann and Hitzler 2010; Böhmann et al. 2018; Funk et al. 2019; Jung et al. 2021). Concepts can be viewed as the building blocks of an ontology and in fact concept fitting also has the support of ontology engineering as a main aim. The techniques needed for concept fitting and ontology fitting are, however, quite different. While it is probably unrealistic to assume that an ontology for an entire domain can be built in a single step from a given set of examples, we believe that small portions of the ontology can be constructed this way, thereby supporting a step-by-step development process by a human engineer. Moreover, in ontology-mediated querying there are applications where a more pragmatic view of an ontology seems appropriate: instead of providing a careful and detailed domain representation, one only wants the ontology to support more complete answers for some given query or a small set of queries (Calvanese et al. 2006; Kharlamov et al. 2017; Sequeda et al. 2019). In such a case, an ontology of rather small size may suffice and deriving it from a collection of examples seems natural, close in spirit to query-by-example.

As ontology languages, we concentrate on the expressive yet fundamental DLs \mathcal{ALC} and \mathcal{ALCT} , and as query languages we consider atomic queries (AQs), conjunctive queries (CQs), full CQs (CQs without quantified variables), and unions of conjunctive queries (UCQs). In addition, we study a fitting problem in which the examples only consist of an ABox and where we seek an ontology that is consistent with the positive examples and inconsistent with the negative ones; this is related, but not identical to both AQ-based and full CQ-based fitting. For all of the resulting combinations, we provide effective characterizations and determine the precise complexity of deciding whether a fitting ontology exists. The algorithms that we use to prove the upper bounds are able to produce explicit fitting ontologies.

For consistency-based fitting and for AQs, our characterizations of fitting existence make use of the connection between

ontology-mediated querying and constraint satisfaction problems (CSPs) established in (Bienvenu et al. 2014). While this connection does not extend to full CQs, the intuitions do and in all three cases our characterizations enable a CONP upper bound, both for \mathcal{ALC} - and \mathcal{ALCI} -ontologies. Corresponding lower bounds are easy to obtain by a reduction from the digraph homomorphism problem. We remark that the complexity is thus much lower than that of the associated query entailment problems, meaning to decide whether $\mathcal{A} \cup \mathcal{O} \models q$ for a given ABox \mathcal{A} , ontology \mathcal{O} , and query q . In fact, the complexity of query entailment is EXPTIME-complete for all cases discussed so far (Baader et al. 2017).

For CQs and UCQs, we give a characterization of fitting existence based on the existence of certain forest models \mathcal{I} . These models are potentially infinite, intuitively because the positive examples (\mathcal{A}, q) act similarly to an existential rule: if we homomorphically find \mathcal{A} in \mathcal{I} , then at the same place we must (in a certain, slightly unusual sense) also find q . Thus the existential quantifiers of q may enforce that every element of \mathcal{I} has a successor, resulting in infinity. As a consequence of this effect, the computational complexity of fitting existence for CQs and UCQs turns out to be much higher than for AQs and full CQs: it is 2EXPTIME-complete both for CQs and UCQs, no matter whether we want to fit an \mathcal{ALC} - or an \mathcal{ALCI} -ontology. For \mathcal{ALCI} , the complexity thus coincides with that of query entailment, which is 2EXPTIME-complete both for CQs and UCQs (Lutz 2008). For \mathcal{ALC} , the complexity of the fitting problem is harder than that of the associated entailment problems, which are both EXPTIME-complete (Lutz 2008). Our upper bounds are obtained by a mosaic procedure. The lower bounds for \mathcal{ALCI} are proved by reduction from query entailment and for \mathcal{ALC} they are proved by reduction from the word problem of exponentially space-bounded alternating Turing machines.

Proofs are provided in the appendix.

Related Work. To the best of our knowledge, the only other study of fitting problems for ontologies is a recent one by (Jung, Hosemann, and Lutz 2025). However, it uses interpretations as examples rather than ABox and queries. Vaguely related are fitting problems for DL concepts. These have been investigated from a practical angle by (Lehmann and Hitzler 2010; Böhmann et al. 2018; Rizzo, Fanizzi, and d’Amato 2020), and from a foundational perspective by (Funk et al. 2019; Jung et al. 2020; Jung et al. 2021; Jung et al. 2022). Other approaches that support the construction of an entire ontology include Angluin’s framework of exact learning by (Konev et al. 2017) and formal concept analysis by (Baader et al. 2007; Baader and Distel 2009; Kriegel 2024). These and related approaches are surveyed by (Ozaki 2020).

2 Preliminaries

2.1 Description Logic

Let N_C , N_R , and N_I be countably infinite sets of *concept names*, *role names*, and *individual names*. An *inverse role* takes the form r^- with r a role name, and a *role* is a role

name or an inverse role. If $r = s^-$ is an inverse role, then we set $r^- = s$. An \mathcal{ALCI} -concept C is built according to

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

where A ranges over concept names and r over roles. As usual, we write \perp as abbreviation for $\neg\top$, $C \sqcup D$ for $\neg(\neg C \sqcap \neg D)$, and $\forall r.C$ for $\neg\exists r.\neg C$. An \mathcal{ALC} -concept is an \mathcal{ALCI} -concept that does not use inverse roles.

An \mathcal{ALCI} -ontology is a finite set of *concept inclusions* (CIs) $C \sqsubseteq D$, where C, D are \mathcal{ALCI} -concepts. \mathcal{ALC} -ontologies are defined likewise. We may write $C \equiv D$ as shorthand for $C \sqsubseteq D$ and $D \sqsubseteq C$. An ABox is a finite set of *concept assertions* $A(a)$ and *role assertions* $r(a, b)$ where A is a concept name, r a role name, and a, b are individual names. We use $\text{ind}(\mathcal{A})$ to denote the set of individual names used in \mathcal{A} .

The semantics of concepts is defined as usual in terms of interpretations $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ with $\Delta^{\mathcal{I}}$ the (non-empty) domain and $\cdot^{\mathcal{I}}$ the *interpretation function*, we refer to (Baader et al. 2017) for full details. An interpretation \mathcal{I} satisfies a CI $C \sqsubseteq D$ if $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$, a concept assertion $A(a)$ if $a \in A^{\mathcal{I}}$, and a role assertion $r(a, b)$ if $(a, b) \in r^{\mathcal{I}}$; we thus make the *standard names assumption*. We say that \mathcal{I} is a *model* of an ontology \mathcal{O} , written $\mathcal{I} \models \mathcal{O}$, if it satisfies all concept inclusions in it, and likewise for ABoxes. An ontology is *satisfiable* if it has a model and an ABox is *consistent* with an ontology \mathcal{O} if \mathcal{A} and \mathcal{O} have a common model.

A *homomorphism* from an interpretation \mathcal{I}_1 to an interpretation \mathcal{I}_2 is a mapping $h: \Delta^{\mathcal{I}_1} \rightarrow \Delta^{\mathcal{I}_2}$ such that $d \in A^{\mathcal{I}_1}$ implies $h(d) \in A^{\mathcal{I}_2}$ and $(d, e) \in r^{\mathcal{I}_1}$ implies $(h(d), h(e)) \in r^{\mathcal{I}_2}$ for all concept names A , role names r , and $d, e \in \Delta^{\mathcal{I}_1}$. We write $\mathcal{I}_1 \rightarrow \mathcal{I}_2$ if there exists a homomorphism from \mathcal{I}_1 to \mathcal{I}_2 and $\mathcal{I}_1 \not\rightarrow \mathcal{I}_2$ otherwise. We will also use homomorphisms from ABoxes to ABoxes and from ABoxes to interpretations. These are defined as expected. In particular, homomorphisms from ABox to ABox need not map individual names to themselves, which would trivialize them.

2.2 Queries

A *conjunctive query* (CQ) takes the form $q = \exists \bar{x} \varphi(\bar{x})$ where \bar{x} is a tuple of variables and φ a conjunction of *atoms* $A(t)$ and $r(t, t')$, with $A \in N_C$, $r \in N_R$, and t, t' variables from \bar{x} or individuals from N_I . With $\text{var}(q)$, we denote the set of variables in \bar{x} . We take the liberty to view q as a set of atoms, writing e.g. $\alpha \in q$ to indicate that α is an atom in q . We may also write $r^-(x, y) \in q$ in place of $r(y, x) \in q$. An *atomic query* (AQ) is a CQ of the simple form $A(a)$, with A a concept name. A CQ is *full* if it does not contain any existentially quantified variables. A *union of conjunctive queries* (UCQ) q is a disjunction of CQs. We refer to each of these classes of queries as a *query language*.

A CQ q gives rise to an interpretation \mathcal{I}_q with $\Delta^{\mathcal{I}_q}$ the set of all variables and individuals in q , $A^{\mathcal{I}_q} = \{t \mid A(t) \in q\}$, and $r^{\mathcal{I}_q} = \{(t, t') \mid r(t, t') \in q\}$ for all $A \in N_C$ and $r \in N_R$. With a homomorphism from a CQ q to an interpretation \mathcal{I} , we mean a homomorphism from \mathcal{I}_q to \mathcal{I} that is the identity on all individual names. If we want to emphasize the latter property, we may speak of a *strong* homomorphism. In contrast, a

weak homomorphism from q to \mathcal{I} , as sometimes used in our proofs, need not be the identity on individual names. For an interpretation \mathcal{I} and a UCQ q , we write $\mathcal{I} \models q$ if there is a (strong) homomorphism h from a CQ in q to \mathcal{I} . For an ABox \mathcal{A} and ontology \mathcal{O} , we write $\mathcal{A} \cup \mathcal{O} \models q$ if $\mathcal{I} \models q$ for all models \mathcal{I} of \mathcal{A} and \mathcal{O} .

Note that all queries introduced above are Boolean, that is, they evaluate to true or false instead of producing answers. For the purposes of this paper, however, this is without loss of generality since we admit individual names in queries.

We use $\|O\|$ to denote the *size* of any syntactic object O such as a concept, an ontology, or a query. It is defined as the length of the encoding of O as a word over some suitable alphabet.

An *ALCC-forest model* \mathcal{I} of an ABox \mathcal{A} is a model of \mathcal{A} such that

1. the directed graph $(\Delta^{\mathcal{I}}, \bigcup_r r^{\mathcal{I}} \setminus (\text{ind}(\mathcal{A}) \times \text{ind}(\mathcal{A})))$ is a forest (a disjoint union of trees) and
2. $r^{\mathcal{I}} \cap (\text{ind}(\mathcal{A}) \times \text{ind}(\mathcal{A})) = \{(a, b) \mid r(a, b) \in \mathcal{A}\}$.

ALCC-forest models are defined likewise, but based on the undirected version of the graph in Point 1. In other words, in *ALCC-forest models* all edges must point away from the roots of the trees while this is not the case for *ALCC-forest models*. With the *degree* of an interpretation, we mean the maximal number of neighbors of any element in its domain.

Lemma 1. *Let $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCI}\}$, \mathcal{O} be an \mathcal{L} -ontology, \mathcal{A} an ABox, and q a UCQ. If $\mathcal{A} \cup \mathcal{O} \not\models q$, then there is an \mathcal{L} -forest model \mathcal{I} of \mathcal{A} and \mathcal{O} of degree at most $\|\mathcal{O}\|$ that satisfies $\mathcal{I} \not\models q$.*

The proof of Lemma 1, which can be found for instance in (Lutz 2008), relies on unraveling, which we shall also use in this article. Let \mathcal{I} be an interpretation and $d \in \Delta^{\mathcal{I}}$. A *path* in \mathcal{I} is a sequence $p = d_1 r_1 \cdots d_{n-1} r_{n-1} d_n$ of domain elements d_i from $\Delta^{\mathcal{I}}$ and role names r_i such that $(d_i, d_{i+1}) \in r_i^{\mathcal{I}}$ for $1 \leq i < n$. We say that the path *starts* at d_1 and use $\text{tail}(p)$ to denote d_n .

The *ALC-unraveling* of \mathcal{I} at d is the interpretation \mathcal{U} defined as follows:

$$\begin{aligned} \Delta^{\mathcal{U}} &= \text{set of all paths in } \mathcal{I} \text{ starting at } d \\ A^{\mathcal{U}} &= \{p \mid \text{tail}(p) \in A^{\mathcal{I}}\} \\ r^{\mathcal{U}} &= \{(p, p') \mid p' = pre \text{ for some } e\}. \end{aligned}$$

The *ALCI-unraveling* of \mathcal{I} at d is defined likewise, with the modification that inverses of roles can also appear in paths and that (p, p') is also included in $r^{\mathcal{U}}$ if $p = p'r^-e$. Note that there is a homomorphism from \mathcal{U} to \mathcal{I} that maps every $p \in \Delta^{\mathcal{U}}$ to $\text{tail}(p)$.

2.3 ABox Examples and the Fitting Problem

Let \mathcal{Q} be a query language such as $\mathcal{Q} = \text{AQ}$ or $\mathcal{Q} = \text{CQ}$. An *ABox- \mathcal{Q} example* is a pair (\mathcal{A}, q) with \mathcal{A} an ABox and q a query from \mathcal{Q} such that all individual names that appear in q are from $\text{ind}(\mathcal{A})$.

By a *collection of labeled examples* we mean a pair $E = (E^+, E^-)$ of finite sets of examples. The examples in E^+ are the *positive examples* and the examples in E^- are the

negative examples. We say that \mathcal{O} *fits* E if $\mathcal{A} \cup \mathcal{O} \models q$ for all $(\mathcal{A}, q) \in E^+$ and $\mathcal{A} \cup \mathcal{O} \not\models q$ for all $(\mathcal{A}, q) \in E^-$. The following example illustrates this central notion.

Example 1. *Consider the collection of labeled ABox-UCQ examples $E = (E^+, E^-)$, where*

$$\begin{aligned} E^+ &= \{ (\{\text{authorOf}(a, b), \text{Publication}(b)\}, \text{Author}(a)), \\ &\quad (\{\text{Reviewer}(a)\}, \exists x \text{reviews}(a, x) \wedge \text{Publication}(x)), \\ &\quad (\{\text{Publication}(a)\}, \text{Confpaper}(a) \vee \text{Jarticle}(a)) \}, \end{aligned}$$

and $E^- = \emptyset$. An *ALC-ontology that fits* (E^+, E^-) is

$$\begin{aligned} \mathcal{O} &= \{ \exists \text{authorOf.Publication} \sqsubseteq \text{Author} \\ &\quad \text{Reviewer} \sqsubseteq \exists \text{reviews.Publication} \\ &\quad \text{Publication} \sqsubseteq \text{Confpaper} \sqcup \text{Jarticle} \}. \end{aligned}$$

There are, however, many other fitting *ALC-ontologies* as well, including as an extreme $\mathcal{O}_{\perp} = \{\top \sqsubseteq \perp\}$ and, say,

$$\mathcal{O}' = \mathcal{O} \cup \{ \text{Author} \sqsubseteq \exists \text{authorOf.Reviewer} \}.$$

We can make both of them non-fitting by adding the negative example

$$(\{\text{Author}(a)\}, \exists x \text{authorOf}(a, x) \wedge \text{Reviewer}(x)).$$

Let \mathcal{L} be an ontology language, such as $\mathcal{L} = \mathcal{ALCI}$, and \mathcal{Q} a query language. Then $(\mathcal{L}, \mathcal{Q})$ -ontology fitting is the problem to decide, given as input a collection of labeled ABox- \mathcal{Q} examples E , whether E admits a fitting \mathcal{L} -ontology. We generally assume that the ABoxes used in E have pairwise disjoint sets of individual names. It is not hard to verify that this is without loss of generality because consistently renaming individual names in a collection of examples has no impact on the existence of a fitting ontology.

There is a natural variation of $(\mathcal{L}, \mathcal{Q})$ -ontology fitting where one additionally requires the fitting ontology to be consistent with all ABoxes that occur in positive examples.¹ We then speak of *consistent $(\mathcal{L}, \mathcal{Q})$ -ontology fitting*. The following observation shows that it suffices to design algorithms for $(\mathcal{L}, \mathcal{Q})$ -ontology fitting as originally introduced.

Proposition 1. *Let \mathcal{L} be any ontology language and $\mathcal{Q} \in \{\text{AQ}, \text{FullCQ}, \text{CQ}, \text{UCQ}\}$. Then there is a polynomial time reduction from consistent $(\mathcal{L}, \mathcal{Q})$ -ontology fitting to $(\mathcal{L}, \mathcal{Q})$ -ontology fitting.*

Proof. We exemplarily treat the case $\mathcal{Q} = \text{AQ}$. The other cases are similar. Let E be a collection of labeled ABox-AQ examples. We extend E to a collection E' by adding, for each positive example $(\mathcal{A}, Q(a)) \in E^+$, a negative example $(\mathcal{A}, X(a))$ where X is a concept name that is not mentioned in E . Then E admits a fitting \mathcal{L} -ontology that is consistent with all ABoxes in positive examples if and only if E' admits a fitting \mathcal{L} -ontology. In fact, any \mathcal{L} -ontology that fits E , does not mention X , and is consistent with all ABoxes in positive examples is also a fitting of E' . Conversely, any ontology that fits E' must be consistent with all ABoxes that occur in positive examples as otherwise one of the additional negative examples would be violated. \square

¹Note that it is implicit already in the original formulation that the fitting ontology must be consistent with all ABoxes that occur in negative examples (\mathcal{A}, q) , as otherwise $\mathcal{A} \cup \mathcal{O} \models q$.

3 Consistency-Based Fitting

We start with a version of ontology fitting that is based on ABox consistency rather than on querying. An example is then simply an ABox, and an ontology \mathcal{O} fits a collection of examples $E = (E^+, E^-)$ if \mathcal{A} is consistent with \mathcal{O} for all $\mathcal{A} \in E^+$ and inconsistent with \mathcal{O} for all $\mathcal{A} \in E^-$. We refer to the induced decision problem as *consistent \mathcal{L} -ontology fitting*.² We believe that it is natural to consider this basic case as a warm-up.

Example 2. Consider the collection of labeled ABox examples $E = (E^+, E^-)$, where

- E^+ contains the ABox $\mathcal{A}_1 = \{r(a_1, a_2)\}$ and
- E^- contains the ABox $\mathcal{A}_2 = \{r(b, b)\}$.

Then $\mathcal{O} = \{\exists r.\exists r.\top \sqsubseteq \perp\}$ is an \mathcal{ALC} -ontology that fits (E^+, E^-) . If we swap E^+ and E^- , then there is no fitting \mathcal{ALC} -ontology or \mathcal{ALCI} -ontology.

We start with a characterization of consistent \mathcal{L} -ontology fitting in terms of homomorphisms.

Theorem 1. Let $E = (E^+, E^-)$ be a collection of labeled ABox examples, $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCI}\}$, and $\mathcal{A}^+ = \biguplus E^+$. Then the following are equivalent:

1. E admits a fitting \mathcal{L} -ontology;
2. $\mathcal{A} \not\vdash \mathcal{A}^+$ for all $\mathcal{A} \in E^-$.

Note that the characterizations for \mathcal{ALC} and \mathcal{ALCI} are identical, and thus a collection of labeled ABox-consistency examples admits a fitting \mathcal{ALC} -ontology if and only if it admits a fitting \mathcal{ALCI} -ontology. It is clear from the proofs that further adding role inclusions, see (Baader et al. 2017), does not increase the separating power either. Adding number restrictions, however, has an impact; see Section 7.

The proof of Theorem 1 makes use of the connection between ontology-mediated querying and constraint satisfaction problems (CSPs) established in (Lutz and Wolter 2012). In particular, for the “2 \Rightarrow 1” direction we use the fact that for every ABox \mathcal{A} , one can construct an ontology \mathcal{O} such that for all ABoxes \mathcal{B} that only use concept and role names from \mathcal{A} , the following holds: $\mathcal{B} \rightarrow \mathcal{A}$ if and only if \mathcal{B} is consistent with \mathcal{O} . We apply this choosing $\mathcal{A} = \mathcal{A}^+$.

We obtain an upper bound for consistent ontology fitting by a straightforward implementation of Point 2 of Theorem 1 and a corresponding lower bound by an easy reduction of the homomorphism problem for directed graphs.

Theorem 2. Let $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCI}\}$. Then consistent \mathcal{L} -ontology fitting is CONP -complete.

It might be worthwhile to point out as a corollary of Theorem 1 that negative examples can be treated independently, in the following sense.

Corollary 1. Let $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCI}\}$ and E be a collection of labeled ABox examples, with $E^- = \{\mathcal{A}_1, \dots, \mathcal{A}_n\}$. Then E admits a fitting \mathcal{L} -ontology if and only if for $1 \leq i \leq n$, the collection of ABox examples $(E^+, \{\mathcal{A}_i\})$ admits a fitting \mathcal{L} -ontology.

²Not to be confused with consistent $(\mathcal{L}, \mathcal{Q})$ -ontology fitting as briefly considered in Proposition 1.

We note that, in related fitting settings such as the one studied in (Funk et al. 2019), statements of this form can often be shown in a very direct way rather than via a characterization. This does not appear to be the case here.

4 Atomic Queries

We consider atomic queries and again present a characterization in terms of homomorphisms. These are now in the other direction, from the positive examples to the negative examples, corresponding to the complementation involved in the well-known reductions from ABox consistency to AQ entailment and vice versa. What is more important, however, is that it does no longer suffice to work directly with the (negative) examples. In fact, the positive examples act like a form of implication on the negative examples, similarly to an existential rule (with atomic unary rule head), and as a result we must first suitably enrich the negative examples.

Let $E = (E^+, E^-)$ be a collection of labeled ABox-AQ examples. A *completion* for E is an ABox \mathcal{C} that extends the ABox $\mathcal{A}^- := \biguplus_{(\mathcal{A}, Q(a)) \in E^-} \mathcal{A}$ by concept assertions $Q(b)$ where $b \in \text{ind}(\mathcal{A}^-)$ and Q a concept name that occurs as an AQ in E^+ .

Theorem 3. Let $E = (E^+, E^-)$ be a collection of labeled ABox-AQ examples and let $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCI}\}$. Then the following are equivalent:

1. E admits a fitting \mathcal{L} -ontology;
2. there is a completion \mathcal{C} for E such that
 - (a) for all $(\mathcal{A}, Q(a)) \in E^+$: if h is a homomorphism from \mathcal{A} to \mathcal{C} , then $Q(h(a)) \in \mathcal{C}$;
 - (b) for all $(\mathcal{A}, Q(a)) \in E^-$: $Q(a) \notin \mathcal{C}$.

The announced behavior of positive examples as an implication is reflected by Point 2a. Note that, as in the consistency case, there is no difference between \mathcal{ALC} and \mathcal{ALCI} . The proof of Theorem 3 is similar to that of Theorem 1. It might be worthwhile to note that Theorem 3 does not suggest a counterpart of Corollary 1. Such a counterpart would speak about single positive examples rather than single negative ones, because of the complementation mentioned above. The following example illustrates that it does not suffice to concentrate on a single positive example (nor a single negative one). Intuitively, this is due to the fact that the ABox \mathcal{A} in Point 2a may be disconnected.

Example 3. Consider the collection of labeled ABox-AQ examples $E = (E^+, E^-)$ with

$$E^+ = \{ (\{A_2(a)\}, A_1(a)), (\{A_3(b), A_4(b')\}, A_2(b)) \}$$

$$E^- = \{ (\{A_3(c)\}, A_1(c)), (\{A_4(d)\}, A_5(d)) \}.$$

E does not admit a fitting \mathcal{ALCI} -ontology, which can be seen by applying Theorem 3: by definition, any completion \mathcal{C} must satisfy $\mathcal{A}^- = \{A_3(c), A_4(d)\} \subseteq \mathcal{C}$. To satisfy Condition 2a of Theorem 3, it must then also satisfy $A_2(c) \in \mathcal{C}$ and $A_1(c) \in \mathcal{C}$. But then \mathcal{C} violates Condition 2b of Theorem 3 for the first negative example. If we drop any of the positive examples, we find completions $\mathcal{C} = \mathcal{A}^-$ and $\mathcal{C} = \mathcal{A}^- \cup \{B(c)\}$, respectively, which satisfy Conditions 2a and 2b. Also dropping any negative example leads to a satisfying completion.

In contrast to the case of consistent \mathcal{L} -ontology fitting, a naive implementation of the characterization given in Theorem 3 only gives a Σ_2^P -upper bound: guess the completion \mathcal{C} for E and then co-guess the homomorphisms in Point (a). In the following, we show how to improve this to CONP. The main observation is that we can do better than guessing \mathcal{C} blindly, by treating positive examples as rules.

Definition 1. Let $E = (E^+, E^-)$ be a collection of labeled ABox-AQ examples, and let $\mathcal{A}^- = \biguplus_{(\mathcal{A}, Q(a)) \in E^-} \mathcal{A}$. A refutation candidate for E is an ABox \mathcal{C} that can be obtained by starting with \mathcal{A}^- and then applying the following rule zero or more times:

(R) if $(\mathcal{A}, Q(a)) \in E^+$ and h is a homomorphism from \mathcal{A} to \mathcal{C} , then set $\mathcal{C} = \mathcal{C} \cup \{Q(h(a))\}$.

Rule (R) can add at most $|E^+| \cdot |\text{ind}(\mathcal{A}^-)|$ (and thus only polynomially many) assertions to \mathcal{A}^- .

Proposition 2. Let $E = (E^+, E^-)$ be a collection of labeled ABox-AQ examples and let $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCI}\}$. Then the following are equivalent:

1. E admits no fitting \mathcal{L} -ontology;
2. there is a refutation candidate \mathcal{C} for E such that $Q(a) \in \mathcal{C}$ for some $(\mathcal{A}, Q(a)) \in E^-$.

Note that Point 1 of Proposition 2 is the complement of Point 1 of Theorem 3, and thus \mathcal{C} has a different role: we may co-guess it, in contrast to the \mathcal{C} in Theorem 3 which needs to be guessed. A close look reveals that we indeed obtain a CONP upper bound.

Theorem 4. Let $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCI}\}$. Then (\mathcal{L}, AQ) -ontology fitting is CONP-complete.

Proof. CONP-hardness can be proved as in the consistency case. It thus remains to argue that the complement of the (\mathcal{L}, AQ) -ontology fitting problem is in NP. By Proposition 2, it suffices to guess an ABox \mathcal{C} with the same individuals as \mathcal{A}^- and to verify that (i) \mathcal{C} is a refutation candidate and (ii) $Q(a) \in \mathcal{C}$ for some $(\mathcal{A}, Q(a)) \in E^-$. To verify (i) we may guess, along with \mathcal{C} , a sequence of positive examples $(\mathcal{A}, Q(a)) \in E^+$ with associated homomorphisms from \mathcal{A} that demonstrate the construction of \mathcal{C} from \mathcal{A}^- by repeated applications of Rule (R). The maximum length of the sequence is $|E^+| \cdot |\text{ind}(\mathcal{A}^-)|$. With the sequence at hand, it is then easy to verify deterministically in polynomial time that \mathcal{C} is a refutation candidate. \square

In view of the close connection between ABox consistency and AQ entailment, one may wonder whether the two fitting problems studied in this and the preceding section are, in some reasonable sense, identical. We may ask whether for every instance E of consistent \mathcal{L} -ontology fitting, there is an instance E' of (\mathcal{L}, AQ) -ontology fitting with the same set of fitting ontologies and vice versa. It turns out that neither is the case. For better readability, in the following we refer to ABox examples as ABox-consistency examples. For \mathcal{L} an ontology language and E a collection of labeled examples, let $O_{E, \mathcal{L}}$ be the set of all \mathcal{L} -ontologies that fit E .

Proposition 3. Let $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCI}\}$. Then

1. there exists a collection of ABox-AQ examples E , such that there is no collection of ABox-consistency examples E' with $O_{E, \mathcal{L}} = O_{E', \mathcal{L}}$;
2. there exists a collection of ABox-consistency examples, such that there is no collection of ABox-AQ examples E' with $O_{E, \mathcal{L}} = O_{E', \mathcal{L}}$.

Proof. For Point 1, consider the collection of ABox-AQ examples $E = (E^+, E^-)$ with $E^+ = \{(\{A(a)\}, B_1(a))\}$ and $E^- = \{(\{A(a)\}, B_2(a))\}$. Note that we use $A(a)$ only to ensure that a occurs in the ABoxes. It is easy to see that for $\mathcal{O}_1 = \{\top \sqsubseteq B_1\}$ and $\mathcal{O}_2 = \{\top \sqsubseteq B_2\}$, \mathcal{O}_1 fits E and \mathcal{O}_2 does not. Furthermore, every ABox is consistent with both \mathcal{O}_1 and \mathcal{O}_2 . Therefore, for every collection of ABox-consistency example E' , either $\{\mathcal{O}_1, \mathcal{O}_2\} \subseteq O_{E', \mathcal{L}}$ or $\mathcal{O}_1 \notin O_{E', \mathcal{L}}$ and $\mathcal{O}_2 \notin O_{E', \mathcal{L}}$.

For Point 2, consider the collection of ABox-consistency examples $E = (E^+, \emptyset)$ with $E^+ = \{\{s(a, b)\}\}$ and let E' be any collection of ABox-AQ examples. If there is a negative example $(\mathcal{A}, A(a))$ in E' for some concept name A , then for $\mathcal{O} = \{\top \sqsubseteq A\}$, \mathcal{O} fits E but \mathcal{O} does not fit E' . If there is no negative example in E' , then for $\mathcal{O}' = \{\top \sqsubseteq \perp\}$, \mathcal{O}' fits E' , but \mathcal{O}' does not fit E . \square

5 Full Conjunctive Queries

We next study the case of full conjunctive queries. Technically, it is closely related to both the AQ-based case and the ABox consistency-based case. However, the potential presence of role atoms in queries brings some technical complications.

We call an example (\mathcal{A}, q) inconsistent if any \mathcal{ALCI} -ontology \mathcal{O} that satisfies $\mathcal{A} \cup \mathcal{O} \models q$ is inconsistent with \mathcal{A} , and consistent otherwise. It is easy to see that any example (\mathcal{A}, q) such that q contains a role atom $r(a, b) \notin \mathcal{A}$ must be inconsistent. In fact, this follows from Lemma 1. Conversely, any example (\mathcal{A}, q) such that q does not contain such a role atom is consistent. This is witnessed by the ontology

$$\mathcal{O} = \{\top \sqsubseteq A \mid A(a) \in q\}.$$

Note that an inconsistent positive example (\mathcal{A}, q) expresses the constraint that \mathcal{A} must be inconsistent with the fitting ontology \mathcal{O} and an inconsistent negative example (\mathcal{A}, q) expresses that \mathcal{A} must be consistent with the fitting ontology \mathcal{O} . In view of this, it is clear that, up to swapping positive and negative examples, FullCQ-based fitting generalizes consistency-based fitting. Moreover, it trivially generalizes AQ-based fitting since every AQ is a full CQ.

Proposition 4. Let $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCI}\}$. Then for every collection of ABox-consistency or ABox-AQ examples E , there is a collection of ABox-FullCQ examples such that $O_{E, \mathcal{L}} = O_{E', \mathcal{L}}$.

For the following development, we would ideally like to get rid of inconsistent examples to achieve simpler characterizations.

There is, however, no obvious way to achieve this for inconsistent positive examples. We can get rid of inconsistent negative examples based on the following observation. Assume that a collection of ABox-FullCQ examples E contains

an inconsistent negative example (\mathcal{A}, q) . We replace it with the negative example $(\mathcal{A}, X(a))$ where X is a fresh concept name and $a \in \text{ind}(\mathcal{A})$ is chosen arbitrarily. The set of fitting \mathcal{ALCT} -ontologies for the resulting set of examples E' remains essentially the same.

Lemma 2. *For $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCT}\}$, there is an \mathcal{L} -ontology that fits E if and only if there is one that fits E' .*

Completions for collections of ABox-FullCQ examples are defined in exact analogy with completions for collections of ABox-AQ examples.

Theorem 5. *Let $E = (E^+, E^-)$ be a collection of labeled ABox-FullCQ examples and let $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCT}\}$. Then the following are equivalent:*

1. E admits a fitting \mathcal{L} -ontology;
2. there is a completion \mathcal{C} for E such that
 - (a) for all consistent $(\mathcal{A}, q) \in E^+$: if h is a homomorphism from \mathcal{A} to \mathcal{C} and $Q(a) \in q$, then $Q(h(a)) \in \mathcal{C}$;
 - (b) for all $(\mathcal{A}, q) \in E^-$: there is a $Q(a) \in q$ such that $Q(a) \notin \mathcal{C}$;
 - (c) for all inconsistent $(\mathcal{A}, q) \in E^+$: there is no homomorphism from \mathcal{A} to \mathcal{C} .

The proof of Theorem 5 uses Theorem 3. Note that, once more, there is no difference between \mathcal{ALC} and \mathcal{ALCT} . As in the case of AQs, our characterization suggests only a Σ_2^P upper bound. However, we can get down to CONP in the same way as for AQs. The following is in exact analogy with Definition 1.

Definition 2. *Let $E = (E^+, E^-)$ be a collection of labeled ABox-FullCQ examples, and let $\mathcal{A}^- = \bigsqcup_{(\mathcal{A}, q) \in E^-} \mathcal{A}$. A refutation candidate for E is an ABox \mathcal{C} that can be obtained by starting with \mathcal{A}^- and then applying the following rule zero or more times:*

(R) *if $(\mathcal{A}, q) \in E^+$ is consistent and h is a homomorphism from \mathcal{A} to \mathcal{C} and $Q(a) \in q$, then set $\mathcal{C} = \mathcal{C} \cup \{Q(h(a))\}$.*

The proof of the following is then analogous to that of Proposition 2. Details are omitted.

Proposition 5. *Let $E = (E^+, E^-)$ be a collection of labeled ABox-AQ examples and let $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCT}\}$. Then the following are equivalent:*

1. E admits no fitting \mathcal{L} -ontology;
2. there is a refutation candidate \mathcal{C} for E such that one of the following conditions is satisfied:
 - (a) there is an $(\mathcal{A}, q) \in E^-$ such that $Q(a) \in \mathcal{C}$ for all $Q(a) \in q$;
 - (b) there is an inconsistent $(\mathcal{A}, q) \in E^+$ and a homomorphism from \mathcal{A} to \mathcal{C} .

And finally, the proof of the following is similar to that of Theorem 4.

Theorem 6. *Let $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCT}\}$. Then $(\mathcal{L}, \text{FullCQ})$ -ontology fitting is CONP-complete.*

6 CQs and UCQs

We now turn to conjunctive queries and UCQs, which constitute the most challenging case. This is due to the fact that, since positive examples act as implications, the presence of existentially quantified variables in the query effectively turns these examples into a form of existential rule. Thus, completions as used for AQs and full CQs are no longer finite.

Throughout this section, we assume that ABoxes in positive examples are never empty. This is mainly to avoid dealing with too many special cases in the technical development. We conjecture that admitting empty ABoxes does not change the obtained results.

6.1 Characterization for \mathcal{ALC} and \mathcal{ALCT}

We start with a characterization for the case of UCQs (and thus also CQs) that is similar in spirit to the one for full CQs given in Theorem 5. The characterization applies to both \mathcal{ALC} and \mathcal{ALCT} in a uniform, though not identical way. As already mentioned, finite completions no longer suffice and we replace them with potentially infinite interpretations. There is another interesting view on this: the fitting ontologies constructed (as part of the proofs) in Sections 3 and 4 do not make existential statements, that is, their sets of models are closed under taking induced subinterpretations. This, however, cannot be achieved for CQs and UCQs. We illustrate this by the following example which also shows that, unlike for AQs and full CQs, there is a difference between fitting \mathcal{ALC} -ontologies and fitting \mathcal{ALCT} -ontologies. Induced subinterpretations are defined in exact analogy with induced substructures in model theory.

Example 4. *Consider the collection of ABox-CQ examples $E = (E^+, E^-)$ where*

$$\begin{aligned} E^+ &= \{(\{A_1(a)\}, \exists x r(x, a) \wedge A_2(x)), \\ &\quad (\{A_2(a)\}, \exists x r(x, a) \wedge A_1(x))\} \\ E^- &= \{(\{A_1(a)\}, B(a)), (\{A_2(a)\}, B(a))\}. \end{aligned}$$

Then there is a fitting \mathcal{ALCT} -ontology:

$$\mathcal{O} = \{A_1 \sqsubseteq \exists r^- . A_2, A_2 \sqsubseteq \exists r^- . A_1\}.$$

The set of models of \mathcal{O} is clearly not closed under taking induced subinterpretations. In fact, this is true for every \mathcal{ALCT} -ontology \mathcal{O}' that fits E since any such \mathcal{O}' must (i) logically imply \mathcal{O} and (ii) be consistent with the ABoxes $\{A(a)\}$ and $\{B(a)\}$, due to the negative examples.

Moreover, it is easy to see that there is no \mathcal{ALC} -ontology \mathcal{O} that fits E . This is due to Lemma 1 and the negative examples, ensuring that any such \mathcal{O} would have to be consistent with the ABoxes in E^+ .

We start with a preliminary. Let $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCT}\}$, let \mathcal{A} be an ABox, \mathcal{I} an interpretation, and h a homomorphism from \mathcal{A} to \mathcal{I} . We define an interpretation $\mathcal{I}_{\mathcal{A}, h, \mathcal{L}}$ as follows. Start with interpretation \mathcal{I}_0 :

$$\begin{aligned} \Delta^{\mathcal{I}_0} &= \text{ind}(\mathcal{A}) \\ A^{\mathcal{I}_0} &= \{a \mid h(a) \in A^{\mathcal{I}}\} \quad \text{for all } A \in \mathbf{N}_{\mathcal{C}} \\ r^{\mathcal{I}_0} &= \{(a, b) \mid r(a, b) \in \mathcal{A}\} \quad \text{for all } r \in \mathbf{N}_{\mathcal{R}}. \end{aligned}$$

Then $\mathcal{I}_{\mathcal{A},h,\mathcal{L}}$ is obtained by taking, for every $a \in \text{ind}(\mathcal{A})$, the \mathcal{L} -unraveling of \mathcal{I} at $h(a)$ and disjointly adding it to \mathcal{I}_0 , identifying the root with a . It can be shown that if \mathcal{I} is a model of some \mathcal{L} -ontology \mathcal{O} , then $\mathcal{I}_{\mathcal{A},h,\mathcal{L}}$ is also a model of \mathcal{O} . Informally, we use $\mathcal{I}_{\mathcal{A},h,\mathcal{L}}$ to ‘undo’ the potential identification of individual names by h , in this way obtaining a forest model of \mathcal{A} .

Theorem 7. *Let $E = (E^+, E^-)$ be a collection of labeled ABox-UCQ examples with $E^- \neq \emptyset$ and let $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCT}\}$. Then the following are equivalent:*

1. *there is an \mathcal{L} -ontology \mathcal{O} that fits E ;*
2. *there is an interpretation \mathcal{I} with degree at most $2^{\|E\|^2}$ such that*
 - (a) *$\mathcal{I} = \bigsqcup_{e \in E^-} \mathcal{I}_e$ where, for each $e = (\mathcal{A}, q) \in E^-$, \mathcal{I}_e is an \mathcal{L} -forest model of \mathcal{A} with $\mathcal{I}_e \not\models q$;*
 - (b) *for all $(\mathcal{A}, q) \in E^+$: if h is a homomorphism from \mathcal{A} to \mathcal{I} , then $\mathcal{I}_{\mathcal{A},h,\mathcal{L}} \models q$.*

The proof of Theorem 7 follows the same intuitions as the proofs of our previous characterizations, but is more technical. One challenge is that, in the “ $2 \Rightarrow 1$ ” direction, we first need to construct from an interpretation \mathcal{I} as in the theorem a suitable finite interpretation that we can then use to identify a fitting \mathcal{L} -ontology. For this we adopt the finite model construction for ontology-mediated querying from (Gogacz, Ibáñez-García, and Murlak 2018).

6.2 Upper Bounds

Our aim is to prove the following.

Theorem 8. *Let $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCT}\}$ and $\mathcal{Q} \in \{\text{CQ}, \text{UCQ}\}$. Then $(\mathcal{L}, \mathcal{Q})$ -ontology fitting is in 2^{EXPTIME} .*

It suffices to prove the theorem for $\mathcal{Q} = \text{UCQ}$. We prove it for \mathcal{ALC} and \mathcal{ALCT} simultaneously. We use the characterization provided by Theorem 7 combined with a mosaic procedure, that is, we attempt to assemble the interpretation \mathcal{I} from Point 2 of Theorem 7 by combining small pieces.

Let $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCT}\}$ and assume that we are given a set of ABox-UCQ examples $E_0 = (E_0^+, E_0^-)$. We will often consider maximally connected components of ABoxes and CQs which, for brevity, we simply call *components*. We wish to work with only connected queries in positive examples. This can be achieved as follows. If $(\mathcal{A}, q) \in E_0^+$ with $q = q_1 \vee \dots \vee q_n$ and q_i has components p_1, \dots, p_k , $k > 1$, then we replace (\mathcal{A}, q) with positive examples $(\mathcal{A}, \hat{q}_1), \dots, (\mathcal{A}, \hat{q}_k)$ where \hat{q}_j is obtained from q by replacing the disjunct q_i with p_j . This leads to an exponential blowup of the number of positive examples, which, however, does not compromise our upper bound because the size of the examples themselves does not increase.

Throughout this section, we shall be concerned with \mathcal{L} -forest models \mathcal{I} of ABoxes \mathcal{A} . We generally assume the following naming convention in such models. All elements of $\Delta^{\mathcal{I}}$ must be of the form aw where $a \in \text{ind}(\mathcal{A})$ and $w \in \mathbb{N}^*$, that is, w is a finite word over the infinite alphabet \mathbb{N} . Moreover, $(d, e) \in r^{\mathcal{I}}$ implies that $d, e \in \text{ind}(\mathcal{A})$ or $e = dc$ or $d = ec$ (if $\mathcal{L} = \mathcal{ALCT}$) where $c \in \mathbb{N}$. If $d, e \in \Delta^{\mathcal{I}}$ and $e = dc$, then we call e a *successor* of d . Note that a

successor may be connected to its predecessor via a role name, an inverse role, or not connected at all. The *depth* of aw is defined as the length of w .

Since mosaics represent ‘local’ pieces of an interpretation, disconnected ABoxes in examples pose a challenge: a homomorphism may map their components into different parts of a forest model that are far away from each other. We thus need some preparation to deal with disconnected ABoxes. For positive examples, one important ingredient is the following observation.

Lemma 3. *Let \mathcal{I} be an interpretation and $(\mathcal{A}, q) \in E^+$ a positive example such that Condition (b) from Theorem 7 is satisfied and each CQ in q is connected. Then there exists a component \mathcal{B} of \mathcal{A} such that: if h is a homomorphism from \mathcal{A} to \mathcal{I} , then $\mathcal{I}_{\mathcal{B},h,\mathcal{L}} \models q$*

Note that Lemma 3 requires the component \mathcal{B} to be uniform across all homomorphisms h . For each $e = (\mathcal{A}, q) \in E^+$, we choose a component $\text{ch}(e)$ of \mathcal{A} . Intuitively, $\text{ch}(e)$ is the component \mathcal{B} from Lemma 3 with \mathcal{I} the interpretation from Point 2 of Theorem 7. Since, however, we do not know \mathcal{I} , ch acts like a guess and our algorithm shall iterate over all possible choice functions ch .

To deal with an example $e = (\mathcal{A}, q) \in E^+$, we shall focus on the component $\text{ch}(e)$ of \mathcal{A} . The other components of \mathcal{A} , however, cannot be ignored. We need to know whether they have a homomorphism to \mathcal{I} , possibly some remote part of it. This is not easily possible from the local perspective of a mosaic, so we again resort to guessing. We choose a set \mathfrak{A} of ABoxes that are a component of the ABox of some positive example. We will take care (locally!) that no ABox in \mathfrak{A} admits a homomorphism to \mathcal{I} . All other components of ABoxes in positive examples may or may not have a homomorphism to \mathcal{I} , we shall simply treat them as if they do. We say that a positive example $(\mathcal{A}, q) \in E^+$ is \mathfrak{A} -enabled if no component of \mathcal{A} is in \mathfrak{A} .

Note that the number of choices for \mathfrak{A} and ch is double exponential (it is not single exponential since we have an exponential number of positive examples, see above).

The queries in negative examples need not be connected. To falsify a non-connected CQ, it clearly suffices to falsify one of its components. We use another choice function to choose these components: for each $e = (\mathcal{A}, p_1 \vee \dots \vee p_k) \in E^-$ and $1 \leq i \leq k$, choose a component $\text{ch}(e, p_i)$ of p_i . There are single exponentially many choices.

Our mosaic procedure tries to assemble the \mathcal{L} -forest model \mathcal{I} starting from a large piece that contains the ABox part of \mathcal{I} as well as the tree parts up to depth $3\|E\|$. The potentially infinite remainder of the trees is then assembled from smaller pieces. We start with defining the large pieces.

Definition 3. *A base candidate for ch and \mathfrak{A} is an interpretation $\mathcal{J} = \bigsqcup_{e \in E^-} \mathcal{I}_e$ that satisfies the following conditions:*

1. *for each $e = (\mathcal{A}, p_1 \vee \dots \vee p_k) \in E^-$, \mathcal{I}_e is an \mathcal{L} -forest model of \mathcal{A} such that $\mathcal{I}_e \not\models \text{ch}(e, p_i)$ for $1 \leq i \leq k$;*
2. *no ABox from \mathfrak{A} has a homomorphism to \mathcal{J} ;*
3. *\mathcal{J} has depth at most $3\|E_0\|$ and degree at most $2^{\|E_0\|^2}$;*

4. for all $e = (\mathcal{A}, q) \in E^+$ that are \mathfrak{A} -enabled: if h is a homomorphism from $\text{ch}(e)$ to \mathcal{J} whose range contains only elements of depth at most $2\|E_0\|$, then $\mathcal{J}_{\text{ch}(e),h,\mathcal{L}} \models q$.

To make sure that there are only finitely many (in fact double exponentially many) base candidates, we assume that (i) \mathcal{J} interprets only concept and role names that occur in E and (ii) if $w \in \Delta^{\mathcal{J}}$ has k successors wc_1, \dots, wc_k , then $\{c_1, \dots, c_k\} = \{1, \dots, k\}$.

We next define the small mosaics. An \mathcal{L} -tree interpretation \mathcal{J} is defined exactly like an \mathcal{L} -forest model, except that all domain elements are of the form $w \in \mathbb{N}^*$, that is, there is no leading individual name. We additionally require the domain $\Delta^{\mathcal{J}}$ to be prefix-closed and call $\varepsilon \in \Delta^{\mathcal{J}}$ the root of \mathcal{J} .

We say that \mathcal{J}' is a subtree of a tree interpretation \mathcal{J} if, for some successor c of ε , \mathcal{J}' is the restriction of \mathcal{J} to all domain elements of the form cw , with $w \in \mathbb{N}^*$.

Definition 4. A mosaic for ch , \mathfrak{A} , and $e = (\mathcal{A}, p_1 \vee \dots \vee p_k) \in E^-$ is an \mathcal{L} -tree interpretation \mathcal{M} that satisfies the following conditions:

1. $\mathcal{M} \not\models \text{ch}(e, p_i)$ for $1 \leq i \leq k$;
2. no ABox from \mathfrak{A} has a homomorphism to \mathcal{M} ;
3. \mathcal{M} has depth at most $3\|E_0\|$ and degree at most $2^{\|E_0\|^2}$;
4. for all $e = (\mathcal{A}, q) \in E^+$ that are \mathfrak{A} -enabled: if h is a homomorphism from $\text{ch}(e)$ to \mathcal{M} whose range contains only elements of depth at least $\|E_0\|$ and at most $2\|E_0\|$, then $\mathcal{M}_{\text{ch}(e),h,\mathcal{L}} \models q$.

Let \mathcal{I} be an \mathcal{L} -forest model and $d \in \Delta^{\mathcal{I}}$. With $\mathcal{I}|_d^\downarrow$, we mean the restriction of \mathcal{I} to all elements of the form dw , with $w \in \mathbb{N}^*$. We say that a mosaic \mathcal{M} glues to d in \mathcal{I} if $\mathcal{I}|_d^\downarrow$ is identical to the interpretation obtained from \mathcal{M} in the following way:

- remove all elements of depth exactly $3\|E_0\|$;
- prefix every domain element with d , that is, every $w \in \Delta^{\mathcal{M}}$ is renamed to dw .

Our algorithm now works as follows. In an outer loop, we iterate over all possible choices for ch and \mathfrak{A} . For each ch and \mathfrak{A} , as well as for each $e \in E^-$ we construct the set $S_{e,0}$ of all mosaics for ch , \mathfrak{A} , and e and then apply an elimination procedure, producing a sequence of sets

$$S_{e,0} \supseteq S_{e,1} \supseteq S_{e,2} \supseteq \dots$$

More precisely, $S_{e,i+1}$ is the subset of mosaics $\mathcal{M} \in S_{e,i}$ that satisfy the following condition:

- (*) for all successors c of ε , there is an $\mathcal{M}' \in S_{e,i}$ that glues to c in \mathcal{M} .

Let S_e be the set of mosaics obtained after stabilization.

We next iterate over all base candidates $\mathcal{J} = \bigcup_{e \in E^-} \mathcal{I}_e$ for ch and \mathcal{A} , for each of them checking whether there is, for every element $d \in \Delta^{\mathcal{I}_e}$ of depth 1, a mosaic $\mathcal{M} \in S_e$ that glues to d in $\Delta^{\mathcal{I}_e}$. If the check succeeds for some ch , \mathfrak{A} and \mathcal{J} , we return ‘fitting exists’. Otherwise, we return ‘no fitting exists’.

Lemma 4. The algorithm returns ‘fitting exists’ if and only if there is an \mathcal{L} -ontology that fits E_0 .

It remains to verify that the algorithm runs in double exponential time. Most importantly, we need an effective way to check Condition 4 of Definitions 3 and 4. This is provided by the subsequent lemma.

Let \mathcal{A} be an ABox and p a CQ. An \mathcal{A} -variation of p is a CQ p' that can be obtained from p by consistently replacing zero or more variables with individual names from $\text{ind}(\mathcal{A})$ and possibly identifying variables. We say that p' is proper if the following conditions are satisfied:

1. if $r(a, b) \in p'$ with $a, b \in \mathbb{N}_1$, then $r(a, b) \in \mathcal{A}$;
2. $\mathcal{I}_{p'}$ is an \mathcal{L} -forest model of $\mathcal{A} \cap p'$.

Further, let \mathcal{I} be an interpretation, h a homomorphism from \mathcal{A} to \mathcal{I} , p' an \mathcal{A} -variation of p and g a weak homomorphism from p' to \mathcal{I} . We say that g is compatible with h if

1. $h(a) = g(a)$ for all individual names a in p' ;
2. for every variable x in p' , there is an $a \in \text{ind}(\mathcal{A})$ such that $g(x)$ is \mathcal{L} -reachable from $h(a)$ in \mathcal{I} .

Here, an element $e \in \Delta^{\mathcal{I}}$ is \mathcal{ALCI} -reachable from $d \in \Delta^{\mathcal{I}}$ if there are $d_0, \dots, d_n \in \Delta^{\mathcal{I}}$ such that $d = d_0$, $d_n = e$, and, for $0 \leq i < n$, $(d_i, d_{i+1}) \in r^{\mathcal{I}}$ for some role r . In \mathcal{ALCI} -reachability, ‘role’ is replaced by ‘role name’.

Lemma 5. Let (\mathcal{A}, q) be an example, \mathcal{I} an interpretation, and h a homomorphism from \mathcal{A} to \mathcal{I} . Then the following are equivalent

1. $\mathcal{I}_{\mathcal{A},h,\mathcal{L}} \models q$
2. there exists a proper \mathcal{A} -variation p' of a CQ p in q and a weak homomorphism from p' to \mathcal{I} that is compatible with h .

The Conditions in Point 2 of Lemma 5 can clearly be checked by brute force in single exponential time, and so can Conditions 1 to 3 of Definitions 3 and 4. Based on Conditions 3 in these definitions, it is thus easy to see that we can produce the set of all base candidates and of all mosaics by a straightforward enumeration in double exponential time. The elimination phase of the algorithm also clearly needs only double exponential time.

6.3 Lower Bounds

Theorem 9. The $(\mathcal{ALCI}, \text{CQ})$ -ontology fitting problem is 2EXPTIME-hard.

To prove Theorem 9, we give a polynomial time reduction from the complement of Boolean CQ entailment in \mathcal{ALCI} , which is 2EXPTIME-hard (Lutz 2007). Assume that we are given \mathcal{A} , \mathcal{O} , and q , and want to decide whether $\mathcal{A} \cup \mathcal{O} \models q$. We construct a collection of labeled ABox-CQ examples (E^+, E^-) such that $\mathcal{A} \cup \mathcal{O} \not\models q$ if and only if there is an \mathcal{ALCI} -ontology \mathcal{O}' that fits (E^+, E^-) .

To keep the reduction simple, we assume \mathcal{O} to be in normal form, meaning that every concept inclusion in \mathcal{O} has one of the following forms: $\top \sqsubseteq A$, $A_1 \sqcap A_2 \sqsubseteq A$, $A \sqsubseteq \exists r.B$, $\exists r.B \sqsubseteq A$, $A \sqsubseteq \neg B$, $\neg B \sqsubseteq A$. It is well-known and easy to see that any \mathcal{ALCI} -ontology \mathcal{O} can be rewritten into an ontology \mathcal{O}' of this form in polynomial time, introducing fresh concept names as needed, such that $\mathcal{A} \cup \mathcal{O} \models q'$ if and only if $\mathcal{A} \cup \mathcal{O}' \models q'$ for all Boolean CQs q' that do not use the fresh concept names.




Figure 1: Arrows denote s -edges.

The reduction uses fresh concept names Real , Choice , $\overline{\text{Choice}}$, and F , a fresh concept name \bar{A} for every concept name A in \mathcal{O} and q , and a fresh role name s . It is helpful to have the characterization in Theorem 7 in mind when reading on.

We use a single negative example to ensure that the interpretation \mathcal{I} from Point 2 of Theorem 7 is a model of \mathcal{A} and makes the concept name F false everywhere:

$$(\mathcal{A} \cup \{\text{Real}(a) \mid a \in \text{ind}(\mathcal{A})\}, \exists x F(x)).$$

We will use a gadget that introduces auxiliary domain elements. To distinguish the domain elements of primary interest from the auxiliary ones, we label the former with the concept name Real .

The interesting part of the reduction is to guarantee that at each element labeled with Real and for each concept name A in \mathcal{O} or q , exactly one of the concept names A and \bar{A} is true. It is easy to express that both concept names cannot be true simultaneously, via the following positive example:

$$(\{A(a), \bar{A}(a)\}, \exists x F(x)).$$

To ensure that at least one of A and \bar{A} is true, we use the announced gadget. We first introduce one successor that satisfies Choice and one that satisfies $\overline{\text{Choice}}$, via the following positive examples:

$$\begin{aligned} &(\{\text{Real}(a)\}, q) \text{ with } q = \exists x (s(a, x) \wedge \text{Choice}) \\ &(\{\text{Real}(a)\}, q) \text{ with } q = \exists x (s(a, x) \wedge \overline{\text{Choice}}). \end{aligned}$$

We then use a positive example (\mathcal{A}^*, q^*) where \mathcal{A}^* and q^* are displayed in Figure 1. To understand the gadget, recall Condition (b) of Theorem 7 and the fact that $\mathcal{I}_{\mathcal{A}^*, h, \mathcal{L}}$ is a forest model of \mathcal{A}^* , for any \mathcal{I} . The variable x in q^* has two distinct individual names as a predecessor, but the only elements in $\mathcal{I}_{\mathcal{A}^*, h, \mathcal{L}}$ with this property are from $\text{ind}(\mathcal{A}^*)$. It follows that any homomorphism that witnesses $\mathcal{I}_{\mathcal{A}^*, h, \mathcal{L}} \models q^*$ as required by Condition (b) of Theorem 7 maps x to c or to d . We transfer the choice back to the original element:

$$\begin{aligned} &(\mathcal{A}, A(a)) \text{ with } \mathcal{A} = \{\text{Real}(a), s(a, b), \text{Choice}(b), A(b)\} \\ &(\mathcal{A}, \bar{A}(a)) \text{ with } \mathcal{A} = \{\text{Real}(a), s(a, b), \overline{\text{Choice}}(b), A(b)\}. \end{aligned}$$

We next include positive examples that encode \mathcal{O} :

$$\begin{aligned} &(\{\text{Real}(a)\}, A(a)) && \text{for every } \top \sqsubseteq A \in \mathcal{O} \\ &(\{A_1(a), A_2(a)\}, A(a)) && \text{for every } A_1 \sqcap A_2 \sqsubseteq A \in \mathcal{O} \\ &(\{A(a)\}, q) \text{ with } q = \exists y (r(a, y) \wedge \text{Real}(y) \wedge B(y)) && \text{for every } A \sqsubseteq \exists r. B \in \mathcal{O} \\ &(\{r(a, b), B(b)\}, A(a)) && \text{for every } \exists r. B \sqsubseteq A \in \mathcal{O} \\ &(\{A(a)\}, \bar{B}(a)) && \text{for every } A \sqsubseteq \neg B \in \mathcal{O} \\ &(\{\bar{B}(a)\}, A(a)) && \text{for every } \neg B \sqsubseteq A \in \mathcal{O}. \end{aligned}$$

Finally, we add a positive example that ensures that F is non-empty if q is made true:

$$(\mathcal{A}_q, \exists x F(x)),$$

where \mathcal{A}_q is q viewed as an ABox, that is, variables become individuals and atoms become assertions. It remains to show the following.

Lemma 6. $\mathcal{A} \cup \mathcal{O} \not\models q$ if and only if there is an \mathcal{ALCC} -ontology that fits (E^+, E^-) .

The reduction used in the proof of Theorem 9 also works for \mathcal{ALC} . But since CQ entailment in \mathcal{ALC} is only EXPTIME-complete, this does not deliver the desired lower bound. We thus resort to a reduction of the word problem for exponentially space-bounded alternating Turing machines (ATMs). Such reductions have been used oftentimes for DL query entailment problems, see e.g. (Lutz 2007; Eiter et al. 2009; Bednarczyk and Rudolph 2022).

Theorem 10. The $(\mathcal{ALC}, \text{CQ})$ -ontology fitting problem is 2EXPTIME-hard.

The crucial step in an ATM reduction of this kind is to ensure that tape cells of two consecutive configurations are labeled in a matching way. This is typically achieved by copying the labeling of each configuration to all successor configurations so that the actual comparison can take place locally. We achieve this with a gadget that is based on the same basic idea as the gadget used in the proof of Theorem 9, but much more intricate.

7 Conclusion

We introduced ontology fitting problems based on ABox-query examples and presented algorithms and complexity results, concentrating on the ontology languages \mathcal{ALC} and \mathcal{ALCC} . We believe that our results can be adapted to cover many common extensions of these. As an illustration, we show in the appendix the following result for the extension \mathcal{ALCCQ} of \mathcal{ALC} with qualified number restrictions. A homomorphism h from an ABox \mathcal{A}_1 to an ABox \mathcal{A}_2 is *locally injective* if $h(b) \neq h(c)$ for all $r(a, b), r(a, c) \in \mathcal{A}_1$.

Theorem 11. Let $E = (E^+, E^-)$ be a collection of labeled ABox examples and $\mathcal{A}^+ = \biguplus E^+$. Then the following are equivalent:

1. E admits a fitting \mathcal{ALCCQ} -ontology;
2. there is no homomorphism from any $A \in E^-$ to \mathcal{A}^+ that is locally injective.

Apart from extensions of \mathcal{ALC} , there are many other natural ontology languages of interest that can be studied in future work, including Horn DLs such as \mathcal{EL} and existential rules. One can also vary the framework in several natural ways and, for instance, consider the case where a signature for the fitting ontology is given as an additional input or where negative examples have a stronger semantics, namely $\mathcal{A} \cup \mathcal{O} \models \neg q$ in place of $\mathcal{A} \cup \mathcal{O} \not\models q$.

Acknowledgements

The third author was supported by DFG project LU 1417/4-1.

References

- Baader, F., and Distel, F. 2009. Exploring finite models in the description logic. In Ferré, S., and Rudolph, S., eds., *Proc. of ICFCA 2009*, volume 5548 of *Lecture Notes in Computer Science*, 146–161. Springer.
- Baader, F.; Ganter, B.; Sertkaya, B.; and Sattler, U. 2007. Completing description logic knowledge bases using formal concept analysis. In Veloso, M. M., ed., *Proc. of IJCAI 2007*, 230–235.
- Baader, F.; Horrocks, I.; Lutz, C.; and Sattler, U. 2017. *An Introduction to Description Logic*. Cambridge University Press.
- Bednarczyk, B., and Rudolph, S. 2022. The price of selfishness: Conjunctive query entailment for ALCSelf is 2EXPTIME-hard. In *36th Proc. of AAAI 2022*, 5495–5502. AAAI Press.
- Bienvenu, M., and Ortiz, M. 2015. Ontology-mediated query answering with data-tractable description logics. In *Proc. of Reasoning Web*, volume 9203 of *LNCS*, 218–307. Springer.
- Bienvenu, M.; ten Cate, B.; Lutz, C.; and Wolter, F. 2014. Ontology-based data access: A study through disjunctive datalog, CSP, and MMSNP. *ACM Trans. Database Syst.* 39(4):33:1–33:44.
- Bourhis, P., and Lutz, C. 2016. Containment in monadic disjunctive datalog, MMSNP, and expressive description logics. In *15th Proc. KR 2016*, 207–216. AAAI Press.
- Bühmann, L.; Lehmann, J.; Westphal, P.; and Bin, S. 2018. DL-Learner structured machine learning on semantic web data. In *Proc. of WWW*, 467–471. ACM.
- Calvanese, D.; Dragone, L.; Nardi, D.; Rosati, R.; and Trisolini, S. 2006. Enterprise modeling and data warehousing in telecom italia. *Inf. Syst.* 31(1):1–32.
- Chandra, A. K.; Kozen, D.; and Stockmeyer, L. J. 1981. Alternation. *J. ACM* 28(1):114–133.
- Eiter, T.; Lutz, C.; Ortiz, M.; and Simkus, M. 2009. Query answering in description logics with transitive roles. In Boutilier, C., ed., *Proc. of IJCAI 2009*, 759–764.
- Funk, M.; Jung, J. C.; Lutz, C.; Pulcini, H.; and Wolter, F. 2019. Learning description logic concepts: When can positive and negative examples be separated? In *Proc. of IJCAI 2019*, 1682–1688. ijcai.org.
- Gogacz, T.; Ibáñez-García, Y. A.; and Murlak, F. 2018. Finite query answering in expressive description logics with transitive roles. In *Proc. of KR 2018*, 369–378. AAAI Press.
- Jacindha, S.; Abishek, G.; and Vasuki, P. 2022. Program synthesis—a survey. In *Computational Intelligence in Machine Learning*, 409–421. Singapore: Springer Nature Singapore.
- Jung, J. C.; Lutz, C.; Pulcini, H.; and Wolter, F. 2020. Logical separability of incomplete data under ontologies. In *Proc. of KR*, 517–528.
- Jung, J. C.; Lutz, C.; Pulcini, H.; and Wolter, F. 2021. Separating data examples by description logic concepts with restricted signatures. In *Proc. of KR*, 390–399.
- Jung, J. C.; Lutz, C.; Pulcini, H.; and Wolter, F. 2022. Logical separability of labeled data examples under ontologies. *Artif. Intell.* 313:103785.
- Jung, J. C.; Hosemann, S.; and Lutz, C. 2025. Fitting ontologies and constraints to relational structures. In *Proc. of KR*.
- Kharlamov, E.; Hovland, D.; Skjæveland, M. G.; Bilidas, D.; Jiménez-Ruiz, E.; Xiao, G.; Soylu, A.; Lanti, D.; Rezk, M.; Zheleznyakov, D.; Giese, M.; Lie, H.; Ioannidis, Y. E.; Kotidis, Y.; Koubarakis, M.; and Waaler, A. 2017. Ontology based data access in Statoil. *J. Web Semant.* 44:3–36.
- Konev, B.; Lutz, C.; Ozaki, A.; and Wolter, F. 2017. Exact learning of lightweight description logic ontologies. *J. Mach. Learn. Res.* 18:201:1–201:63.
- Kriegel, F. 2024. Efficient axiomatization of OWL 2 EL ontologies from data by means of formal concept analysis. In Wooldridge, M. J.; Dy, J. G.; and Natarajan, S., eds., *Proc. of AAAI 2024*, 10597–10606. AAAI Press.
- Lehmann, J., and Hitzler, P. 2010. Concept learning in description logics using refinement operators. *Mach. Learn.* 78(1-2):203–250.
- Li, H.; Chan, C.-Y.; and Maier, D. 2015. Query from examples: An iterative, data-driven approach to query construction. *Proc. VLDB Endow.* 8(13):2158–2169.
- Lutz, C., and Wolter, F. 2012. Non-uniform data complexity of query answering in description logics. In *Proc. KR 2012*. AAAI Press.
- Lutz, C.; Piro, R.; and Wolter, F. 2011. Description logic TBoxes: Model-theoretic characterizations and rewritability. In Walsh, T., ed., *22nd Proc. IJCAI 2011*, 983–988. IJCAI/AAAI.
- Lutz, C. 2007. Inverse roles make conjunctive queries hard. In *Proc. of DL 2007*, volume 250 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Lutz, C. 2008. Two upper bounds for conjunctive query answering in SHIQ. In *Proc. of DL 2008*, volume 353 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Ozaki, A. 2020. Learning description logic ontologies: Five approaches. where do they stand? *Künstliche Intell.* 34(3):317–327.
- Rizzo, G.; Fanizzi, N.; and d’Amato, C. 2020. Class expression induction as concept space exploration: From DL-Foil to DL-Focl. *Future Gener. Comput. Syst.* 108:256–272.
- Sequeda, J. F.; Briggs, W. J.; Miranker, D. P.; and Heide-man, W. P. 2019. A pay-as-you-go methodology to design and build enterprise knowledge graphs from relational databases. In *Proc. of ISWC*, volume 11779 of *LNCS*, 526–545. Springer.
- Shalev-Shwartz, S., and Ben-David, S. 2014. *Understanding Machine Learning - From Theory to Algorithms*. Cambridge University Press.
- Xiao, G.; Calvanese, D.; Kontchakov, R.; Lembo, D.; Poggi, A.; Rosati, R.; and Zakharyashev, M. 2018. Ontology-based data access: A survey. In *Proc. of IJCAI*, 5511–5519. ijcai.org.

A Proofs for Section 3

To prove Theorem 1, we make use of the connection between ontology-mediated querying and constraint satisfaction problems (CSPs) established in (Lutz and Wolter 2012). Theorems 20 and 22 of that paper state the following. A *signature* is a set Σ of concept and role names. For any syntactic object O such as an ontology, an ABox, or a collection of examples, we use $\text{sig}(O)$, we denote the set of concept and role names used in O . With a Σ -ABox, we mean an ABox \mathcal{A} with $\text{sig}(\mathcal{A}) \subseteq \Sigma$.

Proposition 6.

1. For every \mathcal{ALCT} -ontology \mathcal{O} and finite signature $\Sigma \supseteq \text{sig}(\mathcal{O})$, there is a Σ -ABox $\mathcal{A}_{\mathcal{O}}$ such that for all Σ -ABoxes \mathcal{A} : \mathcal{A} is consistent with \mathcal{O} if and only if $\mathcal{A} \rightarrow \mathcal{A}_{\mathcal{O}}$;
2. For every ABox \mathcal{A} and finite signature $\Sigma \supseteq \text{sig}(\mathcal{A})$, there is an \mathcal{ALC} -ontology $\mathcal{O}_{\mathcal{A},\Sigma}$ such that for all Σ -ABoxes \mathcal{B} : $\mathcal{B} \rightarrow \mathcal{A}$ if and only if \mathcal{B} is consistent with $\mathcal{O}_{\mathcal{A},\Sigma}$.

For the reader's information, we recall the construction of the ontology $\mathcal{O}_{\mathcal{A},\Sigma}$ from Point 2. This is of interest because, in the proof of Theorem 1, that ontology will be used as the fitting ontology (if there is one). We introduce a fresh concept name V_a for every $a \in \text{ind}(\mathcal{A})$ and then define $\mathcal{O}_{\mathcal{A},\Sigma}$ to contain the following concept inclusions.

$$\begin{aligned} \top &\sqsubseteq \bigsqcup_{a \in \text{ind}(\mathcal{A})} V_a \\ V_a \sqcap V_b &\sqsubseteq \perp \quad a, b \in \text{ind}(\mathcal{A}) \text{ with } a \neq b \\ V_a &\sqsubseteq \neg A \quad a \in \text{ind}(\mathcal{A}), A \in \Sigma, A(a) \notin \mathcal{A} \\ V_a &\sqsubseteq \forall r. \neg V_b \quad a, b \in \text{ind}(\mathcal{A}), r \in \Sigma, r(a, b) \notin \mathcal{A}. \end{aligned}$$

By construction of $\mathcal{O}_{\mathcal{A},\Sigma}$, a common model \mathcal{I} of a Σ -ABox \mathcal{B} and $\mathcal{O}_{\mathcal{A},\Sigma}$ gives rise to a homomorphism h from \mathcal{B} to \mathcal{A} by setting $h(b) = a$ if $b \in V_a$ for all $a \in \text{ind}(\mathcal{B})$. This is well-defined since \mathcal{I} is a model of $\mathcal{O}_{\mathcal{A},\Sigma}$, and thus the sets $V_a^{\mathcal{I}}$ form a partition of $\Delta^{\mathcal{I}}$.

Theorem 1. Let $E = (E^+, E^-)$ be a collection of labeled ABox examples, $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCT}\}$, and $\mathcal{A}^+ = \bigsqcup E^+$. Then the following are equivalent:

1. E admits a fitting \mathcal{L} -ontology;
2. $\mathcal{A} \not\rightarrow \mathcal{A}^+$ for all $\mathcal{A} \in E^-$.

Proof. “1 \Rightarrow 2”. Let \mathcal{O} be an \mathcal{ALCT} -ontology that fits E . Assume, to the contrary of what we have to show, that $\mathcal{A} \rightarrow \mathcal{A}^+$ for some $\mathcal{A} \in E^-$. Since \mathcal{O} fits E , every $\mathcal{A} \in E^+$ is consistent with \mathcal{O} . By taking the disjoint union of models which witness this, we obtain a model of \mathcal{A}^+ and \mathcal{O} and thus \mathcal{A}^+ is consistent with \mathcal{O} . Therefore, $\mathcal{A}^+ \rightarrow \mathcal{A}_{\mathcal{O}}$ where $\mathcal{A}_{\mathcal{O}}$ is the Σ -ABox from Point 1 of Proposition 6, for $\Sigma = \text{sig}(\mathcal{O}) \cup \text{sig}(E)$. By composition of homomorphisms, we obtain $\mathcal{A} \rightarrow \mathcal{A}_{\mathcal{O}}$. By choice of $\mathcal{A}_{\mathcal{O}}$, this implies that \mathcal{A} is consistent with \mathcal{O} . This contradicts the fact that \mathcal{O} fits E .

“2 \Rightarrow 1”. Assume that $\mathcal{A} \not\rightarrow \mathcal{A}^+$ for all $\mathcal{A} \in E^-$. We argue that the \mathcal{ALC} -ontology $\mathcal{O}_{\mathcal{A}^+,\Sigma}$ from Point 2 of Theorem 6 fits E , for $\Sigma = \text{sig}(E)$. First, let $\mathcal{A} \in E^+$ be a positive example. Clearly, $\mathcal{A} \rightarrow \mathcal{A}^+$. Thus, \mathcal{A} is consistent with $\mathcal{O}_{\mathcal{A}^+,\Sigma}$ by choice of $\mathcal{O}_{\mathcal{A}^+,\Sigma}$. Now let $\mathcal{A} \in E^-$ be a negative example. By assumption, $\mathcal{A} \not\rightarrow \mathcal{A}^+$ and thus \mathcal{A} is not consistent with $\mathcal{O}_{\mathcal{A}^+,\Sigma}$. \square

Theorem 2. Let $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCT}\}$. Then consistent \mathcal{L} -ontology fitting is CONP -complete.

Proof. Theorem 1 places the complement of the consistent \mathcal{L} -ontology fitting problem in NP: construct \mathcal{A}^+ in polynomial time, guess an $\mathcal{A} \in E^-$ and a homomorphism from \mathcal{A} to \mathcal{A}^+ . The lower bound is by a straightforward reduction of the homomorphism problem between directed graphs G_1 and G_2 . Simply use G_1 , viewed as an ABox in the obvious way, as the only negative example and G_2 as the only positive example; then invoke Theorem 1. \square

B Proofs for Section 4

We again use a result that connects fitting to CSPs, in the style of Proposition 6, but for ABoxes in place of consistency (Bourhis and Lutz 2016).

Proposition 7.

1. For every \mathcal{ALCT} -ontology \mathcal{O} and finite signature $\Sigma \supseteq \text{sig}(\mathcal{O})$, there is a Σ -ABox $\mathcal{A}_{\mathcal{O}}$ such that for all Σ -ABoxes \mathcal{A} , $a \in \text{ind}(\mathcal{A})$, and ABoxes \mathcal{Q} : $\mathcal{A} \cup \mathcal{O} \not\models \mathcal{Q}(a)$ if and only if there is a homomorphism h from \mathcal{A} to $\mathcal{A}_{\mathcal{O}}$ with $Q(h(a)) \notin \mathcal{A}_{\mathcal{O}}$;
2. For every ABox \mathcal{A} and finite signature $\Sigma \supseteq \text{sig}(\mathcal{A})$, there is an \mathcal{ALC} -ontology $\mathcal{O}_{\mathcal{A},\Sigma}$ such that for all Σ -ABoxes \mathcal{B} , $b \in \text{ind}(\mathcal{B})$, and concept names Q : there is a homomorphism h from \mathcal{B} to \mathcal{A} with $Q(h(b)) \notin \mathcal{A}$ if and only if $\mathcal{B} \cup \mathcal{O}_{\mathcal{A},\Sigma} \not\models Q(b)$.

Theorem 3. Let $E = (E^+, E^-)$ be a collection of labeled ABox-AQ examples and let $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCT}\}$. Then the following are equivalent:

1. E admits a fitting \mathcal{L} -ontology;
2. there is a completion \mathcal{C} for E such that
 - (a) for all $(\mathcal{A}, Q(a)) \in E^+$: if h is a homomorphism from \mathcal{A} to \mathcal{C} , then $Q(h(a)) \in \mathcal{C}$;
 - (b) for all $(\mathcal{A}, Q(a)) \in E^-$: $Q(a) \notin \mathcal{C}$.

Proof. “1 \Rightarrow 2”. Assume that there is an \mathcal{ALCT} -ontology \mathcal{O} that fits E . Take any $(\mathcal{A}, Q(a)) \in E^-$. Then $\mathcal{A} \cup \mathcal{O} \not\models Q(a)$ and thus \mathcal{A} and \mathcal{O} have a common model \mathcal{I} with $a \notin Q^{\mathcal{I}}$. By taking the disjoint union of these models for all $(\mathcal{A}, Q(a)) \in E^-$, we obtain a model \mathcal{I} of \mathcal{A}^- and \mathcal{O} such that $a \notin Q^{\mathcal{I}}$ for all $(\mathcal{A}, Q(a)) \in E^-$. We define an ABox \mathcal{C} as follows:

$$\begin{aligned} \mathcal{C} &:= \mathcal{A}^- \cup \\ &\{Q(b) \mid (\mathcal{A}, Q(a)) \in E^+, b \in \text{ind}(\mathcal{A}^-), b \in Q^{\mathcal{I}}\}. \end{aligned}$$

It is straightforward to verify that \mathcal{C} is a completion for E that satisfies Condition (b). It remains to show that it satisfies also Condition (a).

Let $(\mathcal{A}, Q(a)) \in E^+$ and let h be a homomorphism from \mathcal{A} to \mathcal{C} . Assume to the contrary of what we have to show that $Q(h(a)) \notin \mathcal{C}$. By definition of \mathcal{C} , this implies $h(a) \notin Q^{\mathcal{I}}$. Thus \mathcal{I} witnesses that $\mathcal{C} \cup \mathcal{O} \not\models Q(h(a))$. Let $\mathcal{A}_{\mathcal{O}}$ be the Σ -ABox from Point 1 of Proposition 7, where $\Sigma = \text{sig}(E) \cup \text{sig}(\mathcal{O})$. Since $\mathcal{C} \cup \mathcal{O} \not\models Q(h(a))$, by choice of $\mathcal{A}_{\mathcal{O}}$ there is a homomorphism g from \mathcal{C} to $\mathcal{A}_{\mathcal{O}}$ with $Q(g(h(a))) \notin \mathcal{A}_{\mathcal{O}}$. Then $f = g \circ h$ is a homomorphism from \mathcal{A} to $\mathcal{A}_{\mathcal{O}}$ such that

$Q(f(a)) \notin \mathcal{A}_O$. By choice of \mathcal{A}_O , this implies $\mathcal{A} \cup \mathcal{O} \not\models Q(a)$, in contradiction to the fact that \mathcal{O} fits E .

“ $2 \Rightarrow 1$ ”. Let \mathcal{C} be a completion for E that satisfies Conditions (a) and (b). Consider the ontology $\mathcal{O}_{\mathcal{C}, \Sigma}$ from Point 2 of Proposition 7, for $\Sigma \supseteq \text{sig}(E)$. We argue that $\mathcal{O}_{\mathcal{C}, \Sigma}$ fits E . First let $(\mathcal{A}, Q(a)) \in E^+$. Then by Condition (a) every homomorphism h from \mathcal{A} to \mathcal{C} satisfies $Q(h(a)) \in \mathcal{C}$. By choice of $\mathcal{O}_{\mathcal{C}, \Sigma}$, this gives $\mathcal{A} \cup \mathcal{O}_{\mathcal{C}, \Sigma} \models Q(a)$, as required. Now let $(\mathcal{A}, Q(a)) \in E^-$. By Condition (b), we have $Q(a) \notin \mathcal{C}$. The identity may thus serve as a homomorphism h from \mathcal{A} to \mathcal{C} with $Q(h(a)) \notin \mathcal{C}$. By choice of $\mathcal{O}_{\mathcal{C}, \Sigma}$, this implies $\mathcal{A} \cup \mathcal{O}_{\mathcal{C}, \Sigma} \not\models Q(a)$, as required. \square

Proposition 2. *Let $E = (E^+, E^-)$ be a collection of labeled ABox-AQ examples and let $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCI}\}$. Then the following are equivalent:*

1. E admits no fitting \mathcal{L} -ontology;
2. there is a refutation candidate \mathcal{C} for E such that $Q(a) \in \mathcal{C}$ for some $(\mathcal{A}, Q(a)) \in E^-$.

Proof. “ $1 \Rightarrow 2$ ”. We prove the contraposition. Assume that there is no refutation candidate \mathcal{C} for E such that $Q(a) \in \mathcal{C}$ for some $(\mathcal{A}, Q(a)) \in E^-$. Let \mathcal{C} be the refutation candidate obtained from \mathcal{A}^- by applying rule R exhaustively. Then $Q(a) \notin \mathcal{C}$ for all $(\mathcal{A}, Q(a)) \in E^-$. Thus, \mathcal{C} is a completion that satisfies Conditions (a) and (b) of Theorem 3. Thus, Theorem 3 implies that E admits a fitting \mathcal{L} -ontology.

“ $2 \Rightarrow 1$ ”. Assume that there is a refutation candidate \mathcal{C} for E such that $Q(a) \in \mathcal{C}$ for some $(\mathcal{A}, Q(a)) \in E^-$. Then every completion \mathcal{C} for E that satisfies Condition (a) of Theorem 3 must fail to satisfy Condition (b). To see this, it suffices to note that the rule R from the definition of refutation candidates is in fact identical to Condition (a). Thus, Point 2 of Theorem 3 fails, and that theorem implies that E does not admit a fitting \mathcal{L} -ontology. \square

C Proofs for Section 5

Lemma 2. *For $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCI}\}$, there is an \mathcal{L} -ontology that fits E if and only if there is one that fits E' .*

Proof. “ \Leftarrow ” Assume that \mathcal{O} is a fitting \mathcal{L} -ontology for E' . We claim that \mathcal{O} is also a fitting \mathcal{L} -ontology for E . Since the sets of positive examples in both collections coincide, it remains to show that $\mathcal{A} \cup \mathcal{O} \not\models q$ for every negative example $(\mathcal{A}, q) \in E$. If (\mathcal{A}, q) is a consistent example, then $(\mathcal{A}, q) \in E'^-$, and therefore $\mathcal{A} \cup \mathcal{O} \not\models q$ by assumption. Now, let $(\mathcal{A}, q) \in E^-$ be an inconsistent example. By construction, $(\mathcal{A}, X(a)) \in E'^+$ and therefore $\mathcal{A} \cup \mathcal{O}$ must be consistent. Using the definition of inconsistent examples, we conclude $\mathcal{A} \cup \mathcal{O} \not\models q$.

“ \Rightarrow ” Assume that \mathcal{O} is a fitting \mathcal{L} -ontology for E . We construct \mathcal{O}' by replacing every occurrence of X in \mathcal{O} with the fresh concept name $Y \in \mathbb{N}_{\mathcal{C}}$. Next, for every interpretation \mathcal{I} , define \mathcal{I}' to be the interpretation obtained from \mathcal{I} by setting $X^{\mathcal{I}'} := \emptyset$ and $Y^{\mathcal{I}'} := X^{\mathcal{I}}$. Observe that for every interpretation \mathcal{J} with $X^{\mathcal{J}} = \emptyset$, there exists an interpretation \mathcal{I} such that $\mathcal{I}' = \mathcal{J}$. Using induction, one can easily show that \mathcal{I} is a model of $\mathcal{A} \cup \mathcal{O}$ if and only if \mathcal{I}' is a model of $\mathcal{A} \cup \mathcal{O}'$, for every ABox \mathcal{A} in which neither X nor Y appears.

Now, let \mathcal{A} be an ABox and q a full conjunctive query, in none of which X or Y occurs.

Claim. $\mathcal{A} \cup \mathcal{O} \models q$ if and only if $\mathcal{A} \cup \mathcal{O}' \models q$.

For the only if direction, assume $\mathcal{A} \cup \mathcal{O}' \not\models q$ and let \mathcal{J} be the respective witness model of $\mathcal{A} \cup \mathcal{O}'$ satisfying $\mathcal{J} \not\models q$. Since X does not occur in \mathcal{O}' , \mathcal{A} or q , we may assume $X^{\mathcal{J}} = \emptyset$. Let \mathcal{I} be the interpretation such that $\mathcal{I}' = \mathcal{J}$. By the earlier argument, \mathcal{I} is a model of $\mathcal{A} \cup \mathcal{O}$. Furthermore, as X and Y do not occur in q , we reason $\mathcal{I} \not\models q$, and hence $\mathcal{A} \cup \mathcal{O} \not\models q$.

To show the if direction, conversely suppose $\mathcal{A} \cup \mathcal{O} \not\models q$. Let \mathcal{I} be the witness model of $\mathcal{A} \cup \mathcal{O}$ with $\mathcal{I} \not\models q$. Then, \mathcal{I}' is a model of $\mathcal{A} \cup \mathcal{O}'$ and again, since neither X nor Y do appear in q , we infer $\mathcal{I}' \not\models q$. Thus, $\mathcal{A} \cup \mathcal{O}' \not\models q$.

We now prove that \mathcal{O}' is a fitting ontology for E' . Let $(\mathcal{A}, q) \in E'^+ = E^+$. Since \mathcal{O} is a fitting ontology for E , $\mathcal{O} \cup \mathcal{A} \models q$ and using the claim above, we derive $\mathcal{O}' \cup \mathcal{A} \models q$. An analogous argument applies to each consistent negative example $(\mathcal{A}, q) \in E^-$. It remains to show that for every inconsistent negative example $(\mathcal{A}, q) \in E^-$, we have $\mathcal{A} \cup \mathcal{O}' \not\models X(a)$. By assumption, there exists a model \mathcal{I} of $\mathcal{A} \cup \mathcal{O}$, and thus \mathcal{I}' is a model of $\mathcal{A} \cup \mathcal{O}'$. Since $X^{\mathcal{I}'} = \emptyset$, we conclude $\mathcal{I}' \not\models q$, and thus \mathcal{O}' is a fitting ontology for E' . \square

Theorem 5. *Let $E = (E^+, E^-)$ be a collection of labeled ABox-FullCQ examples and let $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCI}\}$. Then the following are equivalent:*

1. E admits a fitting \mathcal{L} -ontology;
2. there is a completion \mathcal{C} for E such that
 - (a) for all consistent $(\mathcal{A}, q) \in E^+$: if h is a homomorphism from \mathcal{A} to \mathcal{C} and $Q(a) \in q$, then $Q(h(a)) \in \mathcal{C}$;
 - (b) for all $(\mathcal{A}, q) \in E^-$: there is a $Q(a) \in q$ such that $Q(a) \notin \mathcal{C}$;
 - (c) for all inconsistent $(\mathcal{A}, q) \in E^+$: there is no homomorphism from \mathcal{A} to \mathcal{C} .

Proof. “ $1 \Rightarrow 2$ ”. Assume that E admits a fitting \mathcal{L} -ontology \mathcal{O} . Define a collection E' of ABox-AQ examples as follows:

- $E'^+ = \{(\mathcal{A}, A(a)) \mid (\mathcal{A}, q) \in E^+ \text{ and } A(a) \in q\}$;
- Consider each $(\mathcal{A}, q) \in E^-$; since no inconsistent negative examples are used (\mathcal{A}, q) is consistent. Thus every role atom $r(a, b) \in q$ is an assertion in \mathcal{A} . Since \mathcal{O} fits the negative example (\mathcal{A}, q) , there must thus be a concept assertion $Q(a) \in q$ such that $\mathcal{A} \cup \mathcal{O} \not\models Q(a)$. Include $(\mathcal{A}, Q(a))$ in E'^- .

It is easy to verify that \mathcal{O} fits E' . Thus, there is a completion \mathcal{C} for E' that satisfies Conditions (a) and (b) from Theorem 3. From the proof of that theorem, we additionally know that \mathcal{C} is consistent with \mathcal{O} . By construction of E' , \mathcal{C} being a completion of E' clearly implies that \mathcal{C} also satisfies Conditions (a) and (b) from Theorem 5. It remains to argue that Condition (c) is also satisfied.

To this end, assume that $(\mathcal{A}, q) \in E^+$ is inconsistent. Then \mathcal{A} is inconsistent with \mathcal{O} . Assume to the contrary of what we have to show that there is a homomorphism h from \mathcal{A} to \mathcal{C} .

Let $\mathcal{A}_\mathcal{O}$ be the Σ -ABox from Point 1 of Proposition 6, for $\Sigma = \text{sig}(\mathcal{O}) \cup \text{sig}(E')$. Since \mathcal{C} is consistent with \mathcal{O} , there is a homomorphism g from \mathcal{C} to $\mathcal{A}_\mathcal{O}$. By composing h and g , we obtain a homomorphism from \mathcal{A} to $\mathcal{A}_\mathcal{O}$. As this implies that \mathcal{A} is consistent with \mathcal{O} , we have obtained a contradiction.

“ $2 \Rightarrow 1$ ”. Assume that there is a completion \mathcal{C} for E such that Conditions (a) to (c) from Theorem 5 are satisfied. Define a collection E' of ABox-AQ examples as follows:

- $E'^+ = \{(\mathcal{A}, A(a)) \mid (\mathcal{A}, q) \in E^+ \text{ and } A(a) \in q\}$;
- Consider each $(\mathcal{A}, q) \in E^-$; by Condition (b) and since all negative examples are consistent, there is a $Q(a) \in q$ such that $Q(a) \notin \mathcal{C}$. Include $(\mathcal{A}, Q(a))$ in E'^- .

It is easy to see that \mathcal{C} is also a completion for E' that satisfies Conditions (a) and (b) from Theorem 3. By that theorem, there is therefore an \mathcal{L} -ontology \mathcal{O} that fits E' . We argue that \mathcal{O} also fits E . Since \mathcal{O} fits all negative examples in E' , it is clear the \mathcal{O} fits all negative examples of E .

Let $(\mathcal{A}, q) \in E^+$ be a consistent positive example. Since \mathcal{O} fits all positive examples in E' , we have $\mathcal{A} \cup \mathcal{O} \models Q(a)$ for all concept atoms $Q(a) \in q$. It thus remains to show that $\mathcal{A} \cup \mathcal{O} \models r(a, b)$ for all role atoms $r(a, b) \in q$, but this is clear since (\mathcal{A}, q) is consistent.

Now let $(\mathcal{A}, q) \in E^+$ be an inconsistent positive example. By Condition (c), there is no homomorphism from \mathcal{A} to \mathcal{C} . From the proof of Theorem 3, we actually know that \mathcal{O} can be chosen as the ontology $\mathcal{O}_{\mathcal{C}, \Sigma}$ from Point 2 of Proposition 7, where $\Sigma = \text{sig}(E') \cup \{Q\}$ for a fresh concept name Q . Now, choose some $b \in \text{ind}(\mathcal{A})$. Since there is no homomorphism from \mathcal{A} to \mathcal{C} , there is also no homomorphism h from \mathcal{A} to \mathcal{C} with $Q(h(b)) \notin \mathcal{C}$. By choice of $\mathcal{O}_{\mathcal{C}, \Sigma}$, this implies $\mathcal{A} \cup \mathcal{O}_{\mathcal{C}, \Sigma} \models Q(b)$. But since Q is a fresh concept name, \mathcal{A} must then be inconsistent with $\mathcal{O}_{\mathcal{C}, \Sigma}$. Thus, $\mathcal{O} = \mathcal{O}_{\mathcal{C}, \Sigma}$ fits (\mathcal{A}, q) . \square

Theorem 6. *Let $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCT}\}$. Then $(\mathcal{L}, \text{FullCQ})$ -ontology fitting is CONP-complete.*

Proof. Again CONP-hardness is inherited from the AQ case and it remains to argue that the complement of the $(\mathcal{L}, \text{FullCQ})$ -ontology fitting problem is in NP. By Proposition 5, it suffices to guess an ABox \mathcal{C} with the same individuals as \mathcal{A}^- and to verify that it is a refutation candidate and that it satisfies at least one of Conditions (a) and (b) in Proposition 5. Verifying that \mathcal{C} is a refutation candidate can be achieved exactly as in the proof of Theorem 4. To verify that \mathcal{C} satisfies at least one of Conditions (a) and (b), we may guess which condition is satisfied and, in the case of Condition (b), also an inconsistent $(\mathcal{A}, q) \in E^+$ and a homomorphism from \mathcal{A} to \mathcal{C} . Verifying that (\mathcal{A}, q) is indeed inconsistent can clearly be done in polynomial time. \square

D Proof of Theorem 7

To prove Theorem 7, we need several tools. One of them is a local variant of \mathcal{ALCT} -bisimulations:

Let $\mathcal{I}_1, \mathcal{I}_2$ be interpretations and $k \geq 0$. A relation S is a k - \mathcal{ALCT} -bisimulation between \mathcal{I}_1 and \mathcal{I}_2 if there is a series of relations $S = S_k \subseteq S_{k-1} \cdots \subseteq S_0$ such that the

following conditions are satisfied for all concept names A , roles r , and $i \geq 1$:

1. if $(d_1, d_2) \in S_0$, then $d_1 \in A^{\mathcal{I}_1}$ if and only if $d_2 \in A^{\mathcal{I}_2}$;
2. if $(d_1, d_2) \in S_i$ and $(d_1, d'_1) \in r^{\mathcal{I}_1}$, then there is a $(d_2, d'_2) \in r^{\mathcal{I}_2}$ with $(d'_1, d'_2) \in S_{i-1}$;
3. if $(d_1, d_2) \in S_i$ and $(d_2, d'_2) \in r^{\mathcal{I}_2}$, then there is a $(d_1, d'_1) \in r^{\mathcal{I}_1}$ with $(d'_1, d'_2) \in S_{i-1}$.

We write $(\mathcal{I}_1, d_1) \sim_{\mathcal{ALCT}, k} (\mathcal{I}_2, d_2)$ if there is a k - \mathcal{ALCT} -bisimulation S between interpretations \mathcal{I}_1 and \mathcal{I}_2 with $(d_1, d_2) \in S$. We also define k - \mathcal{ALC} -bisimulations as the variant of k - \mathcal{ALCT} -bisimulations where r ranges only over role names and write $(\mathcal{I}_1, d_1) \sim_{\mathcal{ALC}, k} (\mathcal{I}_2, d_2)$ if there is a k - \mathcal{ALCT} -bisimulation S between interpretations \mathcal{I}_1 and \mathcal{I}_2 with $(d_1, d_2) \in S$.

The following lemma is standard:

Lemma 7. *Let $\mathcal{J}, \mathcal{J}'$ be interpretations, $e \in \Delta^{\mathcal{J}}, e' \in \Delta^{\mathcal{J}'}$, $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCT}\}$, \mathcal{I} an \mathcal{L} -forest model of $\mathcal{A} = \emptyset$ consisting of a single tree with root d and depth at most n . If there is a homomorphism h from \mathcal{I} to \mathcal{J} with $h(d) = e$ and $(\mathcal{J}, e) \sim_{\mathcal{L}, n} (\mathcal{J}', e')$, then there is also a homomorphism h' from \mathcal{I} to \mathcal{J}' with $h'(d) = e'$.*

Note that if \mathcal{J}' is the \mathcal{L} -unraveling of \mathcal{J} at d , then $(\mathcal{J}, e) \sim_{\mathcal{L}, n} (\mathcal{J}', e)$ for any n .

The theorem below is a direct consequence of the model construction used in the proof of Theorem 1 in (Gogacz, Ibáñez-García, and Murlak 2018).

Theorem 12. *Let \mathcal{A} be an ABox, \mathcal{I} an \mathcal{ALCT} -forest model of \mathcal{A} , and $n \geq 1$. Then, there exists a finite model \mathcal{J} of \mathcal{A} such that for all CQs with at most n variables: for every homomorphism h from q to \mathcal{J} , there is a homomorphism h' from q to \mathcal{I} such that*

$$(\mathcal{I}, h(x)) \sim_{\mathcal{ALCT}, n} (\mathcal{J}, h'(x))$$

for all $x \in \text{var}(q)$.

Moreover, we need to restrict the degree of interpretations used in Point 2 of Theorem 7 to be bounded by an exponential and introduce a stronger version of Condition (b) of Theorem 7 that, informally, ensures that matches of disconnected or Boolean CQs are local.

For an \mathcal{L} -forest models \mathcal{I} of some ABox \mathcal{A} , let the *depth* of an element d be the length of the shortest path from an individual $a \in \mathcal{A}$ to d , or ∞ if no such path exists. For $n \geq 0$, let $\mathcal{I}|_n$ be the restriction of \mathcal{I} to elements of at most depth n .

Lemma 8. *Let $E = (E^+, E^-)$ be a collection of labeled ABox-UCQ examples. If there is an interpretation \mathcal{I} that satisfies Conditions (a) and (b) of Theorem 7, then there exists an interpretation \mathcal{I}' with degree at most $2^{\|E\|^2}$ that satisfies Condition (a), (b) and the following variant of Condition (b):*

(b*) for all $(\mathcal{A}, q) \in E^+$, if h is a homomorphism from \mathcal{A} to \mathcal{I}' , then $\mathcal{I}'_{\mathcal{A}, h, \mathcal{L}} \models q$.

Proof. We modify \mathcal{I} as follows to obtain \mathcal{I}' . Let u be a role name that does not occur in E , and let $n = \|E\|$.

First we modify \mathcal{I} to ensure that Condition (b*) holds. As \mathcal{I} satisfies Condition (b), for each $(\mathcal{A}, q) \in E^+$ and

homomorphism h from \mathcal{A} to \mathcal{I} , there must be a CQ p in q such that $\mathcal{I}_{\mathcal{A},h,\mathcal{L}} \models p$. Let g be a homomorphism from p to $\mathcal{I}_{\mathcal{A},h,\mathcal{L}}$, and h' be the extension of h to a homomorphism from $\mathcal{I}_{\mathcal{A},h,\mathcal{L}}$ to \mathcal{I} . For each connected component p' of p that does not contain an individual from \mathcal{A} , pick a variable $x \in \text{var}(p')$, and an individual $a \in \text{ind}(\mathcal{A})$ and extend \mathcal{I} with $(h'(a), h'(g(x))) \in u^{\mathcal{I}}$.

After this modification, \mathcal{I} satisfies Condition (b*), but is no longer an \mathcal{L} -forest model, and has potentially infinite degree. Then, by unraveling \mathcal{I} , we obtain an \mathcal{L} -forest model \mathcal{J} that satisfies Condition (a) and (b*).

We obtain the desired interpretation \mathcal{I}' as a restriction of \mathcal{J} to a subset of its domain. We start by setting \mathcal{I}' to be the restriction of \mathcal{J} to the individuals of $\bigcup_{(\mathcal{A},q) \in E^-} \mathcal{A}$. Thus, at the start the degree of \mathcal{I}' is bounded by $\|E^-\|$ and \mathcal{I}' satisfies Condition (a).

We then extend \mathcal{I}' by exhaustively applying the following rule for all $(\mathcal{A}, q) \in E^+$:

- (*) if h is a homomorphism from \mathcal{A} to \mathcal{I}' and $\mathcal{I}'_{\mathcal{A},h,\mathcal{L}}|_n \not\models q$, then choose a CQ p of q such that $\mathcal{J}_{\mathcal{A},h,\mathcal{L}}|_n \models p$. Now, from all homomorphisms from p to $\mathcal{J}_{\mathcal{A},h,\mathcal{L}}|_n$ select a homomorphism g' for which $\text{im}(g') \cap \Delta^{\mathcal{I}'_{\mathcal{A},h,\mathcal{L}}}$ is maximal. For each component p' of p , let \mathcal{J}' be the restriction of $\mathcal{J}_{\mathcal{A},h,\mathcal{L}}$ to the minimal set of elements that contains the image of p' under g' and is connected to some individual $a \in \mathcal{A}$. Extend \mathcal{I}' with the image of \mathcal{J}' under h' , where h' is the natural extension of h to $\mathcal{J}_{\mathcal{A},h,\mathcal{L}}$.

Exhaustive and fair application of (*), such that every homomorphism h is eventually processed, ensures that \mathcal{I}' satisfies Condition (b*). Since it always remains a restriction of \mathcal{J} to a subset of its domain, it also satisfies Condition (a).

We now show that the degree of any element $d \in \Delta^{\mathcal{I}'}$ is bounded by $2^{\|E\|^2}$. For this, we make use of Lemma 5 and the definition of \mathcal{A} -variations from Section 6.2. It is easy to see that this bound holds for the initial interpretation \mathcal{I}' before any rule application. Now consider a successor $d' \in \Delta^{\mathcal{I}'}$ of d , introduced by a rule application for the positive example $(\mathcal{A}, q) \in E^+$ and homomorphism h from \mathcal{A} to \mathcal{I}' . Let p be the CQ of q and g' the homomorphism from p to $\mathcal{J}_{\mathcal{A},h,\mathcal{L}}|_n$ selected by the rule. By the same argument as in the proof of Lemma 5, there exists a proper \mathcal{A} -variation p' of p along with a weak homomorphism $\hat{h} = h' \circ g'$ from p' to \mathcal{I}' that is compatible with h , where h' denotes the extension of h to $\mathcal{I}'_{\mathcal{A},h,\mathcal{L}}$. Since d' lies in the image of \hat{h} , for each $t \in \hat{h}^{-1}(d')$, there is a subtree of p' rooted at t that maps into the subtree of \mathcal{I}' rooted at d' . Due to the choice of g' , no other successor of d allows a mapping from all of these subtrees. Consequently, for every such subtree of an \mathcal{A} -variation of p , at most one distinct successor is introduced over the course of the entire algorithm.

This restricts the number of successors introduced per positive example $(\mathcal{A}, q) \in E^+$ to

$$(\|(\mathcal{A}, q)\| + 1)^{\|q\|} = 2^{\|q\| \cdot \log(\|(\mathcal{A}, q)\| + 1)}.$$

Summing this over all positive examples and the initial interpretation yields the overall bound $2^{\|E\|^2}$ on the degree of any element in $\Delta^{\mathcal{I}'}$. \square

Theorem 7. Let $E = (E^+, E^-)$ be a collection of labeled ABox-UCQ examples with $E^- \neq \emptyset$ and let $\mathcal{L} \in \{\mathcal{ALCC}, \mathcal{ALCIT}\}$. Then the following are equivalent:

1. there is an \mathcal{L} -ontology \mathcal{O} that fits E ;
2. there is an interpretation \mathcal{I} with degree at most $2^{\|E\|^2}$ such that
 - (a) $\mathcal{I} = \biguplus_{e \in E^-} \mathcal{I}_e$ where, for each $e = (\mathcal{A}, q) \in E^-$, \mathcal{I}_e is an \mathcal{L} -forest model of \mathcal{A} with $\mathcal{I}_e \not\models q$;
 - (b) for all $(\mathcal{A}, q) \in E^+$: if h is a homomorphism from \mathcal{A} to \mathcal{I} , then $\mathcal{I}_{\mathcal{A},h,\mathcal{L}} \models q$.

Proof. “1 \Rightarrow 2”. Let \mathcal{O} be an \mathcal{L} -ontology that fits (E^+, E^-) . Then $\mathcal{A} \cup \mathcal{O} \not\models q$ for all $(\mathcal{A}, q) \in E^-$. By Lemma 1, we can thus choose for every $e = (\mathcal{A}, q) \in E^-$ an \mathcal{L} -forest model of \mathcal{A} and \mathcal{O} as \mathcal{I}_e such that $\mathcal{I}_e \not\models q$ and degree at most $\|\mathcal{O}\|$. By construction, $\mathcal{I} := \biguplus_{e \in E^-} \mathcal{I}_e$ satisfies Condition (a). Note that $\mathcal{I} := \biguplus_{e \in E^-} \mathcal{I}_e$ is a well-defined interpretation (i.e., it has a non-empty domain) due to our assumption that $E^- \neq \emptyset$.

Next, we show that \mathcal{I} also satisfies Condition (b). For this let $(\mathcal{A}, q) \in E^+$ and let h be a homomorphism from \mathcal{A} to \mathcal{I} . Observe that \mathcal{I} is a model of \mathcal{O} since each \mathcal{I}_e is. This implies that $\mathcal{I}_{\mathcal{A},h,\mathcal{L}}$ is not only a model of \mathcal{A} , but also a model of \mathcal{O} . Since \mathcal{O} fits E^+ , it must be that $\mathcal{A} \cup \mathcal{O} \models q$, and thus $\mathcal{I}_{\mathcal{A},h,\mathcal{L}} \models q$, as required.

It then follows from Lemma 8 that there also exists an interpretation \mathcal{I}' with degree at most $2^{\|E\|^2}$ that satisfies Conditions (a) and (b).

“2 \Rightarrow 1”. Let $(\mathcal{I}_e)_{e \in E^-}$ be interpretations such that $\mathcal{I} = \biguplus_{e \in E^-} \mathcal{I}_e$ satisfies Conditions (a) and (b). Without loss of generality we assume that \mathcal{I} interprets as non-empty only the concept and role names in E . We proceed as follows: first, we convert \mathcal{I} into a finite model \mathcal{J} that also satisfies Conditions (a) and (b). Then, we use \mathcal{J} to construct the desired ontology \mathcal{O} that fits (E^+, E^-) . By Lemma 8, we can assume that \mathcal{I} satisfies Condition (b*).

We now apply Theorem 12 to each \mathcal{I}_e using $n = \|E\|$ to obtain a finite model \mathcal{J}_e .

Claim. $\mathcal{J} = \biguplus_{e \in E^-} \mathcal{J}_e$ satisfies the following variations of Conditions (a) and (b) from Theorem 7:

- (a') for each $e = (\mathcal{A}, q) \in E^-$: \mathcal{J}_e is a model of \mathcal{A} and $\mathcal{J}_e \not\models q$;
- (b') for all $(\mathcal{A}, q) \in E^+$: if h is a homomorphism from \mathcal{A} to \mathcal{J} , then $\mathcal{J}_{\mathcal{A},h,\mathcal{L}} \models q$.

Condition (a') follows directly from Theorem 12 and the choice of n .

For Condition (b'), let h be a homomorphism from \mathcal{A} to \mathcal{J} . As the individual parts \mathcal{J}_e of \mathcal{J} are disjoint, consider connected components of \mathcal{A} individually. If h maps a connected component of \mathcal{A} to \mathcal{J}_e , then (by viewing the component as a Boolean CQ) Theorem 12 implies that there is a homomorphism from this component to \mathcal{I}_e . By combining these homomorphisms we obtain a homomorphism h' from the entirety of \mathcal{A} to \mathcal{I} .

Since \mathcal{I} satisfies Condition (b*), there thus must be a homomorphism g from a CQ p in q to $\mathcal{I}_{\mathcal{A},h',\mathcal{L}}|_n$. Let \mathcal{I}^* be the minimal connected restriction of $\mathcal{I}_{\mathcal{A},h',\mathcal{L}}|_n$ that contains

the image of g . Note that \mathcal{I}^* is an \mathcal{L} -forest model and that the depth of \mathcal{I}^* is at most n . We now argue that there exists a homomorphism from \mathcal{I}^* to $\mathcal{J}_{\mathcal{A},h,\mathcal{L}}$ that is the identity on $\text{ind}(\mathcal{A})$.

For this, we first verify that the identity is a homomorphism from \mathcal{I}^* restricted to domain $\text{ind}(\mathcal{A})$ to $\mathcal{J}_{\mathcal{A},h,\mathcal{L}}$. For this, consider an $a \in \text{ind}(\mathcal{A})$ with $a \in A^{\mathcal{I}^*}$. By definition of $\mathcal{I}_{\mathcal{A},h',\mathcal{L}}$, $h'(a) \in A^{\mathcal{I}_e}$ for some component \mathcal{I}_e of \mathcal{I} . Theorem 12 implies that $(\mathcal{I}_e, h'(a)) \sim_{\mathcal{A}\mathcal{L}\mathcal{C}\mathcal{I},n} (\mathcal{J}_e, h(a))$. Therefore, $h(a) \in A^{\mathcal{J}_e}$ and thus $a \in A^{\mathcal{J}_{\mathcal{A},h,\mathcal{L}}}$.

Now consider a tree-shaped component $\mathcal{I}^{*'}$ of \mathcal{I}^* that is rooted at some $a \in \text{ind}(\mathcal{A})$. Note that $|\Delta^{\mathcal{I}^{*'}}| \leq n$. Thus, again using the fact that $(\mathcal{I}_e, h'(a)) \sim_{\mathcal{A}\mathcal{L}\mathcal{C}\mathcal{I},n} (\mathcal{J}_e, h(a))$, for e such that $h(a) \in \Delta^{\mathcal{J}_e}$, we can conclude, using Lemma 7, that there is a homomorphism from $\mathcal{I}^{*'}$ to \mathcal{J}_e that maps a to $h(a)$, and therefore there is also a homomorphism from $\mathcal{I}^{*'}$ to $\mathcal{J}_{\mathcal{A},h,\mathcal{L}}$ that maps $\mathcal{I}^{*'}$ to the \mathcal{L} -unraveling of \mathcal{J}_e at $h(a)$ that is attached to a in $\mathcal{J}_{\mathcal{A},h,\mathcal{L}}$.

Thus, there is a homomorphism from the entirety of \mathcal{I}^* to $\mathcal{J}_{\mathcal{A},h,\mathcal{L}}$ that is the identity on $\text{ind}(\mathcal{A})$, and by composition of homomorphisms, $\mathcal{J}_{\mathcal{A},h,\mathcal{L}} \models q$, as required.

From \mathcal{J} we now construct an ontology \mathcal{O} that fits the positive and negative examples. The ontology \mathcal{O} uses fresh concept names V_d for each $d \in \Delta^{\mathcal{J}}$ and is constructed as follows:

$$\begin{aligned} \mathcal{O} &= \{ \top \sqsubseteq \bigsqcup_{d \in \Delta^{\mathcal{J}}} V_d \\ V_d \sqcap V_e &\sqsubseteq \perp && \text{for } d, e \in \Delta^{\mathcal{J}} \text{ with } d \neq e \\ V_d &\sqsubseteq A && \text{for } d \in A^{\mathcal{J}} \\ V_d &\sqsubseteq \neg A && \text{for } d \in \Delta^{\mathcal{J}} \setminus A^{\mathcal{J}} \\ V_d &\sqsubseteq \exists r. V_e && \text{for } (d, e) \in r^{\mathcal{J}} \\ V_d &\sqsubseteq \neg \exists r. V_e && \text{for } (d, e) \in (\Delta^{\mathcal{J}} \times \Delta^{\mathcal{J}}) \setminus r^{\mathcal{J}} \end{aligned}$$

where A ranges over concept names that occur in E or have a non-empty extension in \mathcal{J} and r ranges over role names that occur in E or have a non-empty extension in \mathcal{J} as well as over inverse roles if $\mathcal{L} = \mathcal{A}\mathcal{L}\mathcal{C}\mathcal{I}$.

By setting $V_d^{\mathcal{J}} = \{d\}$ for each $d \in \Delta^{\mathcal{J}}$, we extend \mathcal{J} to a model of \mathcal{O} .

Observe that in any model \mathcal{I} of \mathcal{O} , for every $e \in \Delta^{\mathcal{I}}$, there is exactly one $d \in \Delta^{\mathcal{J}}$ such that $e \in V_d^{\mathcal{I}}$. Thus, the relation

$$S = \{(e, d) \in \Delta^{\mathcal{I}} \times \Delta^{\mathcal{J}} \mid e \in V_d^{\mathcal{I}}\}$$

is a function from left to right. Furthermore, using the concept inclusions in \mathcal{O} , one can show that S is an \mathcal{L} -bisimulation and induces (due to its functionality) a homomorphism from \mathcal{I} to \mathcal{J} , if \mathcal{I} only interprets concept and role names that occur in \mathcal{O} as non-empty. Using this property, we now show that \mathcal{O} fits the examples.

First consider a negative example $e = (\mathcal{A}, q) \in E^-$. The component \mathcal{J}_e of \mathcal{J} is a model of \mathcal{O} and \mathcal{A} , and from Condition (a') it follows that $\mathcal{J}_e \not\models q$, as required.

Now consider a positive example $(\mathcal{A}, q) \in E^+$. If there are no common models of \mathcal{A} and \mathcal{O} , then $\mathcal{A} \cup \mathcal{O} \models q$ and we are done. Otherwise, let $\mathcal{I}_{\mathcal{A}}$ be a model of \mathcal{A} and \mathcal{O} . Without

loss of generality assume that $\mathcal{I}_{\mathcal{A}}$ interprets only concept and role names that do occur in \mathcal{O} as non-empty. Note that all concept and role names used in \mathcal{A} must also be used in \mathcal{O} by construction.

By the properties of \mathcal{O} 's models, the relation S , as defined above, is an \mathcal{L} -bisimulation between $\mathcal{I}_{\mathcal{A}}$ and \mathcal{J} , and S directly induces a homomorphism from $\mathcal{I}_{\mathcal{A}}$ to \mathcal{J} . Since $\mathcal{I}_{\mathcal{A}}$ is also a model of \mathcal{A} , the identity is a homomorphism from \mathcal{A} to $\mathcal{I}_{\mathcal{A}}$ and therefore by composition there is a homomorphism h from \mathcal{A} to \mathcal{J} .

As \mathcal{J} satisfies Condition (b'), the existence of h implies that $\mathcal{J}_{\mathcal{A},h,\mathcal{L}} \models q$. Let g thus be a homomorphism from a CQ p in q to $\mathcal{J}_{\mathcal{A},h,\mathcal{L}}$.

We now aim to show that there is homomorphism g' from $\mathcal{J}_{\mathcal{A},h,\mathcal{L}}$ to $\mathcal{I}_{\mathcal{A}}$ that is the identity on $\text{ind}(\mathcal{A})$. The composition of g and g' is then a homomorphism that witnesses $\mathcal{I}_{\mathcal{A}} \models q$ as required.

For this, observe that by construction of $\mathcal{J}_{\mathcal{A},h,\mathcal{L}}$, there is an \mathcal{L} -bisimulation S' between $\mathcal{J}_{\mathcal{A},h,\mathcal{L}}$ and \mathcal{J} with $(a, h(a)) \in S'$. By composing S and S' , we obtain an \mathcal{L} -bisimulation \hat{S} between $\mathcal{J}_{\mathcal{A},h,\mathcal{L}}$ and $\mathcal{I}_{\mathcal{A}}$. Note that since S agrees with h , $(a, a) \in \hat{S}$ for all $a \in \text{ind}(\mathcal{A})$.

We now define g' in two steps. First set $g'(a) = a$ for all $a \in \text{ind}(\mathcal{A})$. Note that by definition of $\mathcal{J}_{\mathcal{A},h,\mathcal{L}}$, g' is a homomorphism from $\mathcal{J}_{\mathcal{A},h,\mathcal{L}}$ restricted to $\text{ind}(\mathcal{A})$ to $\mathcal{I}_{\mathcal{A}}$. Then, let $a \in \text{ind}(\mathcal{A})$ and consider the \mathcal{L} -unraveling of \mathcal{J} that was attached to a in the definition of $\mathcal{J}_{\mathcal{A},h,\mathcal{L}}$. Since $S \circ S'$ is an \mathcal{L} -bisimulation with $(a, a) \in S \circ S'$, we can extend g' to this entire \mathcal{L} -unraveling of \mathcal{J} attached to a .

This completes the definition of g' . \square

E Proofs for Section 6.2

Lemma 3. *Let \mathcal{I} be an interpretation and $(\mathcal{A}, q) \in E^+$ a positive example such that Condition (b) from Theorem 7 is satisfied and each CQ in q is connected. Then there exists a component \mathcal{B} of \mathcal{A} such that: if h is a homomorphism from \mathcal{A} to \mathcal{I} , then $\mathcal{I}_{\mathcal{B},h,\mathcal{L}} \models q$*

Proof. We start with an easy observation: if h is a homomorphism from \mathcal{A} to \mathcal{I} , then there is a component \mathcal{B} of \mathcal{A} such that $\mathcal{I}_{\mathcal{B},h,\mathcal{L}} \models q$. This is due to the fact that Condition (b) from Theorem 7 is satisfied and q is connected. This does not yet give the lemma since the lemma requires us to find a component \mathcal{B} of \mathcal{A} that works uniformly for all h .

Assume to the contrary of what we have to show that for every component \mathcal{B} of \mathcal{A} there is a homomorphism $h_{\mathcal{B}}$ from \mathcal{A} to \mathcal{I} such that $\mathcal{I}_{\mathcal{B},h_{\mathcal{B}},\mathcal{L}} \not\models q$. Let $h_{\mathcal{B}}^-$ be the restriction of $h_{\mathcal{B}}$ to $\text{ind}(\mathcal{B})$ and let h be the union of the homomorphisms $h_{\mathcal{B}}^-$, for all components \mathcal{B} of \mathcal{A} . Then h is a homomorphism from \mathcal{A} to \mathcal{I} and $\mathcal{I}_{\mathcal{B},h,\mathcal{L}} \not\models q$ for all components \mathcal{B} of \mathcal{A} . This, however, contradicts our initial observation. \square

Lemma 4. *The algorithm returns ‘fitting exists’ if and only if there is an \mathcal{L} -ontology that fits E_0 .*

Proof. We first show that if the algorithm returns ‘fitting exists’, then there is an interpretation \mathcal{I} that satisfies Conditions (a) and (b) of Theorem 7 for E .

If the algorithm returns ‘fitting exists’, then there is a choice of \mathfrak{A} and ch , a base candidate $\mathcal{J} = \bigcup_{e \in E^-} \mathcal{I}_e$ for \mathfrak{A} and ch , as well as for each $e \in E^-$ a set S_e of mosaics. We use these to construct the desired interpretation \mathcal{I} .

For this, start with $\mathcal{I} = \mathcal{J} = \bigcup_{e \in E^-} \mathcal{I}_e$, and for every $e \in E^-$ consider \mathcal{I}_e . Since the algorithm returned ‘fitting exists’ there must be, for every element d of depth 1 in \mathcal{I}_e a mosaic $\mathcal{M} \in S_e$ that glues to d . Then extend \mathcal{I}_e by gluing \mathcal{M} to d , by which we mean taking the (non-disjoint) union of \mathcal{I}_e and \mathcal{M} in which every domain element of \mathcal{M} is prefixed with d . If there are multiple candidates for \mathcal{M} , choose arbitrarily. Now continue this extension process for all elements d of depth $i \geq 2$ in a breadth-first manner. Each such element d must originally occur on depth 1 of some mosaic $\mathcal{M} \in S_e$. By Condition (*) on the set S_e , there must be a $\mathcal{M}' \in S_e$ that glues to d . Extend \mathcal{I}_e by gluing \mathcal{M}' to d .

Next, we verify that the resulting interpretation $\mathcal{I} = \bigcup_{e \in E^-} \mathcal{I}_e$ satisfies Conditions (a) and (b) of Theorem 7.

Let $e = (\mathcal{A}, p_1 \vee \dots \vee p_k) \in E^-$. By construction of \mathcal{I} , \mathcal{I}_e is an \mathcal{L} -forest model of \mathcal{A} , and the degree of \mathcal{I} is bound by $2^{\|E_0\|^2}$. Assume for contradiction that $\mathcal{I}_e \models q$. Then, there must be a CQ p_i in q such that $\mathcal{I}_e \models p_i$, which implies that $\mathcal{I}_e \models \text{ch}(e, p_i)$, as $\text{ch}(e, p_i)$ is a connected component of p_i . By definition, $\text{ch}(e, p_i)$ is of size at most $\|E_0\|$ and thus must match either entirely into the base candidate \mathcal{J} of \mathcal{I} or entirely into a mosaic \mathcal{M} . However, both definitions demand in Condition 1 that the component $\text{ch}(e, p_i)$ has no match. Thus, $\mathcal{I}_e \not\models q$ and Condition (a) must hold for \mathcal{I} .

Now let $e = (\mathcal{A}, q) \in E^+$ and h be a homomorphism from \mathcal{A} to \mathcal{I} . The existence of h implies that e is \mathfrak{A} -enabled. As $\text{ch}(e)$ is a component of \mathcal{A} , h is also a homomorphism from $\text{ch}(e)$ to \mathcal{I} . By definition, $\text{ch}(e)$ has size at most $\|E_0\|$. Thus, the range of h restricted to $\text{ch}(e)$ contains either only elements of depth at most $2\|E_0\|$, or only elements of depth between $\|E_0\| + k$ and depth $2\|E_0\| + k$ for some $k \geq 1$. In the first case, the range of h restricted to $\text{ch}(e)$ lies entirely within the base candidate \mathcal{J} of \mathcal{I} . By Condition 4 of base candidates, thus $\mathcal{J}_{\text{ch}(e), h, \mathcal{L}} \models q$ and $\mathcal{I}_{\text{ch}(e), h, \mathcal{L}} \models q$, as required. In the second case, the range of h restricted to $\text{ch}(e)$ lies entirely within depth $\|E_0\|$ and $2\|E_0\|$ of some mosaic \mathcal{M} . Condition 4 of mosaics then implies that $\mathcal{M}_{\text{ch}(e), h, \mathcal{L}} \models q$, and thus $\mathcal{I}_{\text{ch}(e), h, \mathcal{L}} \models q$, as required.

Now we show that if there is an interpretation \mathcal{I} that satisfies Conditions (a) and (b) of Theorem 7, then the algorithm returns ‘fitting exists’.

Let $\mathcal{I} = \bigcup_{e \in E^-} \mathcal{I}_e$ be an interpretation as in Theorem 7 for E_0 . Most importantly, \mathcal{I} has degree at most $2^{\|E_0\|^2}$. By Lemma 8, we may assume that \mathcal{I} also satisfies the Condition (b*). By Lemma 3, for each positive example $e = (\mathcal{A}, q) \in E^+$, there must be a component \mathcal{B} of \mathcal{A} such that if h is a homomorphism from \mathcal{A} to \mathcal{I} , then $\mathcal{I}_{\mathcal{B}, h, \mathcal{L}} \models q$. Choose $\text{ch}(e)$ for all $e \in E^+$ accordingly.

Next, consider all components \mathcal{B} of ABoxes \mathcal{A} that occur in positive examples, and let \mathfrak{A} be the set of all \mathcal{B} that do not have homomorphisms to \mathcal{I} .

Additionally, since \mathcal{I} satisfies Condition (a), $\mathcal{I}_e \not\models q$ for each negative example $e = (\mathcal{A}, p_1 \vee \dots \vee p_k) \in E^-$. Thus,

for each p_i , there must be a choice for $\text{ch}(e, p_i)$ such that $\mathcal{I}_e \not\models \text{ch}(e, p_i)$.

Let \mathcal{J} now be the restriction of \mathcal{I} up to depth $3\|E_0\|$. We show that \mathcal{J} is a base candidate for ch and \mathfrak{A} . Condition 1 of base candidates is satisfied by choice of $\text{ch}(e, p_i)$ and the fact that \mathcal{I} satisfies Condition (a). Condition 2 of base candidates is satisfied by choice of \mathfrak{A} . Condition 3 is satisfied by definition of \mathcal{J} . For Condition 4, consider a positive example $e = (\mathcal{A}, q) \in E^+$ that is \mathfrak{A} -enabled. If h is a homomorphism from $\text{ch}(e)$ to \mathcal{I} whose range contains only element of depth at most $2\|E_0\|$, then by choice of $\text{ch}(e)$, $\mathcal{I}_{\text{ch}(e), h, \mathcal{L}} \models q$. Thus, as \mathcal{I} also satisfies Condition (b*) and the fact that the range of h contains only elements of depth at most $2\|E_0\|$, $\mathcal{J}_{\text{ch}(e), h, \mathcal{L}} \models q$.

Now for the mosaics. Let $e \in E^-$. For each $d \in \Delta^{\mathcal{I}_e}$, let \mathcal{M}_d be the restriction of \mathcal{I}_e to all elements of the form dw with $|w| < 3\|E_0\|$.

We argue that for every for $d \in \Delta^{\mathcal{I}_e}$ with depth at least 1, \mathcal{M}_d is a mosaic for \mathfrak{A} , ch and $e = (\mathcal{A}, p_1 \vee \dots \vee p_k) \in E^-$. Condition 1 of mosaics is satisfied by choice of $\text{ch}(e, p_i)$ and the fact that \mathcal{I} satisfies Condition (a). Condition 2 of mosaics is satisfied by choice of \mathfrak{A} . Condition 3 is satisfied by definition of \mathcal{M}_d . Condition 4 holds by the same argument as in the base candidate case.

Now consider the set

$$S_e = \{\mathcal{M}_d \mid d \in \Delta^{\mathcal{I}_e}, \text{depth of } d \geq 1\}.$$

We claim that S_e is stable, i.e., that for every $\mathcal{M}_d \in S_e$ and every successor c of d , there is an $\mathcal{M}'_c \in S_e$ such that \mathcal{M}'_c glues to c in \mathcal{M}_d . In fact this is easy to see, as $\mathcal{M}_c \in S_e$.

Thus, the algorithm returns ‘fitting exists’. \square

Lemma 5. *Let (\mathcal{A}, q) be an example, \mathcal{I} an interpretation, and h a homomorphism from \mathcal{A} to \mathcal{I} . Then the following are equivalent*

1. $\mathcal{I}_{\mathcal{A}, h, \mathcal{L}} \models q$
2. *there exists a proper \mathcal{A} -variation p' of a CQ p in q and a weak homomorphism from p' to \mathcal{I} that is compatible with h .*

Proof. “1 \Rightarrow 2”. Assume that $\mathcal{I}_{\mathcal{A}, h, \mathcal{L}} \models q$ and let g be a strong homomorphism from a CQ p in q to $\mathcal{I}_{\mathcal{A}, h, \mathcal{L}}$. We use g to identify an \mathcal{A} -variation p' by replacing a variable x with a if $g(x) = a \in \text{ind}(\mathcal{A})$ and identifying variables x and y if $g(x) = g(y) \notin \text{ind}(\mathcal{A})$. Since $\mathcal{I}_{\mathcal{A}, h, \mathcal{L}} \models q$ and $\mathcal{I}_{\mathcal{A}, h, \mathcal{L}}$ is an \mathcal{L} -forest model of \mathcal{A} , this \mathcal{A} -variation is proper. Moreover, g is also a homomorphism from p' to $\mathcal{I}_{\mathcal{A}, h, \mathcal{L}}$. From the construction of $\mathcal{I}_{\mathcal{A}, h, \mathcal{L}}$ it follows that there is also a homomorphism h' from $\mathcal{I}_{\mathcal{A}, h, \mathcal{L}}$ to \mathcal{I} that extends h . The composition $h' \circ g$ is then a weak homomorphism from p' to \mathcal{I} . By construction, $h' \circ g$ satisfies Condition 1 of being compatible with h . Since every element in $\mathcal{I}_{\mathcal{A}, h, \mathcal{L}}$ is \mathcal{L} -reachable from some individual in $\text{ind}(\mathcal{A})$ in $\mathcal{I}_{\mathcal{A}, h, \mathcal{L}}$, $h' \circ g$ also satisfies Condition 2.

“2 \Rightarrow 1”. Assume that there exists a proper \mathcal{A} -variation p' of p and an accompanying weak homomorphism g from p' to \mathcal{I} that is compatible with h . To prove that $\mathcal{I}_{\mathcal{A}, h, \mathcal{L}} \models q$, it clearly suffices to show that there is a strong homomorphism h' from p' to $\mathcal{I}_{\mathcal{A}, h, \mathcal{L}}$. To assemble h' , we start with the

identity mapping on the set I of individuals in p' . Since p' is proper, this is a strong homomorphism from p' restricted to I to $\mathcal{I}_{A,h,\mathcal{L}}$.

Next, take any maximally connected component q_c of $p' \setminus \mathcal{A}$. Clearly, q_c contains at most one individual name and \mathcal{I}_{q_c} is an \mathcal{L} -tree.

We distinguish two cases:

- q_c contains an individual, say a . Clearly g is a homomorphism from q_c to \mathcal{I} with $g(a) = h(a)$. As \mathcal{I}_{q_c} is an \mathcal{L} -tree, then by Lemma 7, there is also a homomorphism from q_c to the \mathcal{L} -unraveling of \mathcal{I} at $h(a)$ that is attached in $\mathcal{I}_{A,h,\mathcal{L}}$ to a . Extend h' accordingly.

- q_c does not contain an individual.

Then, let $x \in \text{var}(q_c)$. As g is compatible with h , there must be an $a \in \text{ind}(\mathcal{A})$ such that $g(x)$ is \mathcal{L} -reachable from $h(a)$ in \mathcal{I} . Thus, as in the previous case, extend h' to map q_c to the \mathcal{L} -unraveling of \mathcal{I} that is attached to $h(a)$ in $\mathcal{I}_{A,h,\mathcal{L}}$.

Overall, we obtain the desired strong homomorphism h' from q' to $\mathcal{I}_{A,h,\mathcal{L}}$. \square

F Proofs for Section 6.3

Lemma 6. $\mathcal{A} \cup \mathcal{O} \not\models q$ if and only if there is an \mathcal{ALCC} -ontology that fits (E^+, E^-) .

Proof. “ \Rightarrow ” If $\mathcal{A} \cup \mathcal{O} \not\models q$, then by Lemma 1 there is an \mathcal{ALCC} -forest model \mathcal{I} of \mathcal{A} and \mathcal{O} of degree at most $\|\mathcal{O}\|$ such that $\mathcal{I} \not\models q$. By construction of E^+ , $\|\mathcal{O}\| \leq \|E^+\| \leq 2\|E\|^2$.

To show that there is a ontology \mathcal{O}' that fits (E^+, E^-) , we show that one can extend \mathcal{I} to an interpretation \mathcal{I}' that satisfies Conditions (a) and (b) of Theorem 7 for the examples (E^+, E^-) .

First, set $\text{Real}^{\mathcal{I}'} = \Delta^{\mathcal{I}'}$. This ensures that \mathcal{I}' is a model of the ABox in the only negative example. Then, for each $d \in \Delta^{\mathcal{I}'}$, introduce fresh individuals d_1 and d_2 with $(d, d_1), (d, d_2) \in s^{\mathcal{I}'}$, $d_1 \in \text{Choice}^{\mathcal{I}'}$ and $d_2 \in \overline{\text{Choice}}^{\mathcal{I}'}$. Additionally, for all concept names A that occur in \mathcal{O} or q , if $d \in A^{\mathcal{I}'}$, add d_1 to $A^{\mathcal{I}'}$ and d_2 to $\overline{A}^{\mathcal{I}'}$; if $d \notin A^{\mathcal{I}'}$, add both d and d_1 to $\overline{A}^{\mathcal{I}'}$ and d_2 to $A^{\mathcal{I}'}$.

Now \mathcal{I}' satisfies Condition (b) for all examples in E^+ . In particular, each element is labeled with exactly one of A, \overline{A} for each concept name A that occurs in \mathcal{O} or q , and there is no homomorphism from \mathcal{A}_q to \mathcal{I}' . Furthermore, by construction $F^{\mathcal{I}'} = \emptyset$, and thus \mathcal{I}' satisfies Condition (a). Therefore, by Theorem 7, there is an \mathcal{ALCC} -ontology that fits (E^+, E^-) .

“ \Leftarrow ” If there is an \mathcal{ALCC} -ontology \mathcal{O}' that fits (E^+, E^-) , then by Theorem 7, there is an interpretation \mathcal{I} that satisfies Conditions (a) and (b). We are then interested in the restriction \mathcal{I}' of \mathcal{I} to the domain $\Delta^{\mathcal{I}'} = \text{Real}^{\mathcal{I}'}$.

We argue that \mathcal{I}' is a model of \mathcal{A} and \mathcal{O} with $\mathcal{I}' \not\models q$. First note that since \mathcal{I}' satisfies Condition (a), it is a model of \mathcal{A} with $F^{\mathcal{I}'} = \emptyset$. By construction of the positive examples that involve Choice and \overline{A} , and the fact that \mathcal{I} satisfies Condition (b), it follows that $(\overline{A})^{\mathcal{I}'} = \Delta^{\mathcal{I}'} \setminus A^{\mathcal{I}'}$ for all concept

names A that occur in \mathcal{O} or q . Then, it follows from the positive examples that encode \mathcal{O} that \mathcal{I}' is a model of \mathcal{O} . For example, consider a concept inclusion $A \sqsubseteq \neg B \in \mathcal{O}$. Then, there exists an example $(\{A(a)\}, \overline{B}(a)) \in E^+$. Let $d \in A^{\mathcal{I}'}$. Then, $d \in \text{Real}^{\mathcal{I}'}$ and there is a homomorphism h from $\{A(a)\}$ to \mathcal{I} with $h(a) = d$, and thus d must also be labeled with $\overline{B}(a)$. Hence, $d \notin B^{\mathcal{I}'}$ as required. The arguments for the other forms of concept inclusions are similar.

As \mathcal{I}' additionally satisfies Condition (b) for the final positive example $(\mathcal{A}_q, \exists x F(x))$, and $F^{\mathcal{I}'} = \emptyset$, there is no homomorphism from \mathcal{A}_q to \mathcal{I}' , and therefore $\mathcal{I}' \not\models q$. \square

Theorem 10. *The (\mathcal{ALCC}, CQ) -ontology fitting problem is 2EXPTIME-hard.*

We prove the theorem using a polynomial time reduction from the word problem of exponentially space-bounded alternating Turing machines (ATMs).

Definition 5 (Alternating Turing Machine). *An alternating Turing machine is a tuple $\mathcal{M} = (Q, \Sigma, q_0, \Delta)$, where*

- Q is a set of states partitioned into a set of existential states Q_{\exists} , a set of universal states Q_{\forall} , an accepting state $\{q_{\text{acc}}\}$ and a rejecting state $\{q_{\text{rej}}\}$,
- $q_0 \in Q_{\exists} \cup Q_{\forall}$ is the starting state,
- Σ is a finite set of symbols, containing a blank symbol \square , and
- $\Delta \subseteq Q \times \Sigma \times Q \times \Sigma \times \{+1, -1\}$ is the transition relation.

To simplify notation, we define $\Delta(q, a) = \{(q', a', M) \mid (q, a, q', a', M) \in \Delta\}$.

Let $\mathcal{M} = (Q, \Sigma, q_0, \Delta)$ be an ATM. A configuration of \mathcal{M} is a word wqw' where $w, w' \in \Sigma^*$ and $q \in Q$ and represents the content of the tape ww' , head position $|w| + 1$, and state q of the Turing machine. Successor configurations of a configuration wqw' are then defined using Δ in the standard way. We assume that a configuration wqw' with $q \in Q_{\exists} \cup Q_{\forall}$ always has at least one successor configuration and that a configuration $wqw' \in \{q_{\text{acc}}, q_{\text{rej}}\}$ has no successor configurations.

A computation of an ATM \mathcal{M} on a word w is sequence of configurations K_0, K_1, \dots such that $K_0 = q_0w$ and K_{i+1} is a successor configuration of K_i for all $i \geq 0$. For our purposes it suffices to consider ATMs that have finite computations on any input. A configuration wqw' without successor configurations is *accepting* if $q = q_{\text{acc}}$. A configuration wqw' with $q \in Q_{\exists}$ is *accepting* if at least one successor configuration is accepting. A configuration wqw' with $q \in Q_{\forall}$ is *accepting* if all successor configurations are accepting. Finally an ATM \mathcal{M} *accept* an input $w \in \Sigma^*$ if the configuration q_0w is accepting. If \mathcal{M} accepts an input w , then this is witnessed by a finite computation tree whose nodes are accepting configurations, the root being the initial configuration q_0w and the leaves being configurations in state q_{acc} or q_{rej} .

There is an exponentially space-bounded ATM $\mathcal{M} = (Q, \Sigma, q_0, \Delta)$ whose word problem is 2EXPTIME-hard, and thus, on input w with $|w| = m$, all reached configurations $w'qw''$ satisfy $|w'w''| < 2^m$ (Chandra, Kozen, and Stockmeyer 1981).

From \mathcal{M} and an input $w \in \Sigma^*$, we now construct a collection $E = (E^+, E^-)$ of ABox-CQ examples such that E admits a fitting \mathcal{ALC} -ontology if and only if \mathcal{M} accepts w . This is based on the characterization of $(\mathcal{ALC}, \text{CQ})$ -ontology fitting in Theorem 7: every interpretation satisfying Conditions (a) and (b) of Theorem 7 will represent an accepting computation of \mathcal{M} on w as a tree of configurations.

For readability, we split the construction of $E = (E^+, E^-)$ into three steps:

1. We define an \mathcal{ALC} -ontology \mathcal{O} such that \mathcal{ALC} -forest models of $\mathcal{O} \cup \{I(a)\}$ represent the structure of accepting computation trees of \mathcal{M} on w , but do not yet ensure that consecutive configurations are labeled in a matching way;
2. from \mathcal{O} , we construct ABox-CQ examples E' such that every interpretation satisfying Conditions (a) and (b) of Theorem 7 satisfies the same properties as models of $\mathcal{O} \cup \{I(a)\}$;
3. we extend E' with examples that ensure that consecutive configurations are labeled in a matching way.

To understand the construction of \mathcal{O} , observe that when started in input w with $|w| = m$, \mathcal{M} only visits configurations of length at most 2^m . We can thus represent configurations of \mathcal{M} at the leaves of binary trees of height m . We include in \mathcal{O} concept inclusions that generate such trees and ensure they are well-labeled.

The roots of such configuration trees are marked with the concept name R . Using the following concept inclusions, we enforce, at each node with label R , the existence of a binary tree of height m and simultaneously assign each leaf its position on the tape using a binary representation based on the concept names B_1, \dots, B_m , where B_1 represents the least significant bit:

$$R \sqsubseteq L_0$$

$$L_i \sqsubseteq \exists r. (L_{i+1} \sqcap B_{i+1}) \sqcap \exists r. (L_{i+1} \sqcap \neg B_{i+1})$$

for i with $0 \leq i < m$, and

$$L_i \sqcap B_j \sqsubseteq \forall r. (L_{i+1} \rightarrow B_j)$$

$$L_i \sqcap \neg B_j \sqsubseteq \forall r. (L_{i+1} \rightarrow \neg B_j),$$

for all i, j with $0 < j \leq i < m$. In the following, we call the elements that are labeled with L_m (tape) cells.

In each cell, we store the current tape symbol and possibly a state $q \in Q$. To simplify the further construction, we also store the tape symbol and state of the same cell with respect to the preceding configuration. We therefore include concept inclusions that add two elements to each cell, one labeled with the concept name M_h , which we will call *h-memory*, and another one labeled with the concept name M_p , which we call *p-memory*, where *h* refers to *here* and *p* refers to *preceding*.

To be able to compare cell positions, at each h- and p-memory element we introduce m additional successors, the i -th successor being labeled with the concept name A_i and storing the i -th bit of the position of the cell using the concept names A and \bar{A} . Each of these successors, in turn, has a single successor to enable the construction of certain gadgets later.

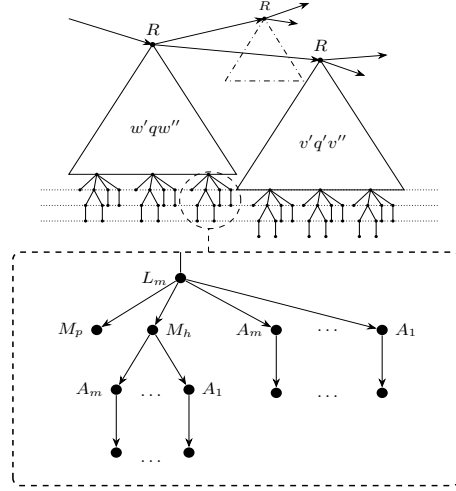


Figure 2: Structure of \mathcal{ALC} -forest models of \mathcal{O} .

For this, we add the following concept inclusions:

$$L_m \sqsubseteq \exists r. M_h \sqcap \exists r. M_p$$

$$L_m \sqsubseteq \exists r. A_i$$

$$M_h \sqsubseteq \exists r. A_i$$

$$L_m \sqcap B_i \sqsubseteq \forall r. (A_i \rightarrow A) \sqcap \forall r. (A_i \rightarrow \bar{A})$$

$$L_m \sqcap \neg B_i \sqsubseteq \forall r. (A_i \rightarrow \bar{A}) \sqcap \forall r. (A_i \rightarrow A)$$

$$\bar{A} \sqsubseteq \neg A$$

$$A_i \sqsubseteq \exists r. \top$$

for all i with $1 \leq i \leq m$.

Next, we add concept inclusions which guarantee that the p-memory and h-memory of every cell are appropriately labeled with tape symbols and states. For this, we use concept names for each $a \in \Sigma$ and $q \in Q \cup \{\text{nil}\} = Q^+$, where nil represents that the head of the ATM is currently in a different cell. For simplicity we directly use the elements of Σ and Q as concept names.

$$M_h \sqcap M_p \sqsubseteq \bigsqcup_{q \in Q^+} (q \sqcap \prod_{q' \in Q^+ \setminus \{q\}} \neg q')$$

$$M_h \sqcap M_p \sqsubseteq \bigsqcup_{a \in \Sigma} (a \sqcap \prod_{a' \in \Sigma \setminus \{a\}} \neg a').$$

For easier comparison, we also label each combination of symbols $a \in \Sigma$ and $q \in Q^+$ with concept name $Z_{a,q}$:

$$M_h \sqcap M_p \sqsubseteq (a \sqcap q) \leftrightarrow Z_{a,q}$$

We store the current and preceding positions of the head of \mathcal{M} in the roots of configuration trees. For this we use a binary encoding via the concept names Q_j and Q'_j , respectively, for $1 \leq j \leq m$. By propagating these concept names to all cells in this configuration, we are able to ensure that only the appropriate cell is labeled with a state. For this, we add the

following concept inclusions:

$$\begin{aligned}
L_i \sqcap Q_j &\sqsubseteq \forall r. (L_{i+1} \rightarrow Q_j) \\
L_i \sqcap \neg Q_j &\sqsubseteq \forall r. (L_{i+1} \rightarrow \neg Q_j) \\
L_i \sqcap Q'_j &\sqsubseteq \forall r. (L_{i+1} \rightarrow Q'_j) \\
L_i \sqcap \neg Q'_j &\sqsubseteq \forall r. (L_{i+1} \rightarrow \neg Q'_j) \\
L_m \sqcap C_{=} &\sqsubseteq \forall r. (M_h \rightarrow \bigsqcup_{q \in Q} q) \\
M_h \sqcap \neg C_{=} &\sqsubseteq \forall r. (M_h \rightarrow \text{nil}),
\end{aligned}$$

for i, j with $0 \leq i < m$ and $1 \leq j \leq m$, $C_{=} := \prod_{i=1}^m (Q_i \leftrightarrow B_i)$, and $C'_{=} := \prod_{i=1}^m (Q'_i \leftrightarrow B_i)$. Further, we need to ensure that configurations have successors that locally comply with the transition relation Δ and respect the structure of accepting ATM computation trees. Hence, we demand the existence of one successor configuration tree if the current configuration is in an existential state, and the existence of all required successor configurations if the current configuration is in a universal state. For this, we store the chosen transition $(q, a, q', a', M) \in \Delta$ in the root of the successor configuration using a concept name $T_{a', q', M}$. Hence, for $q \in Q_{\exists}$ and $a \in \Sigma$, we include

$$R \sqcap \exists r^{m+1} (M_h \sqcap q \sqcap a) \sqsubseteq \bigsqcup_{(q', a', M) \in \Delta(q, a)} \exists r. (R \sqcap T_{q', a', M}).$$

And for $q \in Q_{\forall}$ and $a \in \Sigma$, we include

$$R \sqcap \exists r^{m+1} (M_h \sqcap q \sqcap a) \sqsubseteq \prod_{(q', a', M) \in \Delta(q, a)} \exists r. (R \sqcap T_{q', a', M}).$$

We propagate the old head position to the next configuration while changing the concept names from Q_i to Q'_i . This enables us to distinguish the old from the new head position. For this, we include the following concept inclusions:

$$\begin{aligned}
Q_i \sqcap R &\sqsubseteq \forall r. (R \rightarrow Q'_i) \quad \text{for all } 1 \leq i \leq m \\
\neg Q'_i \sqcap R &\sqsubseteq \forall r. (R \rightarrow \neg Q'_i) \quad \text{for all } 1 \leq i \leq m
\end{aligned}$$

The new head position is then calculated depending on the respective transition node $T_{a', a', M}$ for $M \in \{+1, -1\}$. This is achieved by the following concept inclusions for $M = +1$:

$$\begin{aligned}
R \sqcap T_{q, a, +1} &\sqsubseteq Q'_1 \leftrightarrow \neg Q_1 \\
R \sqcap T_{q, a, +1} \sqcap (Q'_i \sqcap \neg Q_i) &\sqsubseteq (Q'_{i+1} \leftrightarrow \neg Q_{i+1}) \\
R \sqcap T_{q, a, +1} \sqcap (\neg Q'_i \sqcup Q_i) &\sqsubseteq (Q'_{i+1} \leftrightarrow Q_{i+1})
\end{aligned}$$

for all i with $1 \leq i < m$, and with the following inclusions for $M = -1$:

$$\begin{aligned}
R \sqcap T_{q, a, -1} &\sqsubseteq Q'_1 \leftrightarrow \neg Q_1 \\
R \sqcap T_{q, a, -1} \sqcap (\neg Q'_i \sqcap Q_i) &\sqsubseteq (Q'_{i+1} \leftrightarrow \neg Q_{i+1}) \\
R \sqcap T_{q, a, -1} \sqcap (Q'_i \sqcup \neg Q_i) &\sqsubseteq (Q'_{i+1} \leftrightarrow Q_{i+1}),
\end{aligned}$$

for all $1 \leq i < m$.

Using the old and new head positions as well as the transition concept $T_{a', q', M}$, we can enforce that the cells of this configuration correctly store the result of the respective transition. Thus, the cell at the old head position should store the

symbol a' , the cell at the new head position should be labeled with q' , and the symbol stored in all other cells should remain the same. The latter can be enforced by enforcing that for each cell, the symbol stored in the p-memory is the same as the one stored in the h-memory. We assume here that the p-memory is set up correctly, which is not yet guaranteed but will be ensured later on in Step 3 of the reduction.

We include the following concept inclusions for all $T_{q', a', M}$ and $a \in \Sigma$

$$\begin{aligned}
T_{q', a', M} &\sqsubseteq \forall r^m. (L_m \rightarrow T_{q', a', M}) \\
L_m \sqcap T_{q', a', M} \sqcap C'_{=} &\sqsubseteq \forall r. (M_h \rightarrow a') \\
L_m \sqcap T_{q', a', M} \sqcap C_{=} &\sqsubseteq \forall r. (M_h \rightarrow q') \\
L_m \sqcap \exists r. (M_p \sqcap a \sqcap \text{nil}) &\sqsubseteq \forall r. (M_h \rightarrow a),
\end{aligned}$$

where $C'_{=} := \prod_{i=1}^m (Q'_i \leftrightarrow B_i)$.

To ensure that the represented run is accepting, we demand that every configuration without successor states must be in the accepting state:

$$R \sqcap \forall r. \neg R \sqsubseteq \bigsqcup_{a \in \Sigma, M \in \{-1, +1\}} T_{a, q_{\text{acc}}, M}$$

The final concept inclusions are used to encode the initial configuration. For this, we will take care that the root of the computation tree is labeled with the concept name I . We then put:

$$\begin{aligned}
I &\sqsubseteq \exists r. R \\
I &\sqsubseteq \forall r^{m+1}. (\text{pos}_B = i \rightarrow \forall r. (M_h \rightarrow a_i))
\end{aligned}$$

for all i with $0 \leq i < m$, and

$$\begin{aligned}
I &\sqsubseteq \forall r^{m+1}. (\text{pos}_B = 0 \rightarrow \forall r. (M_h \rightarrow q_0)) \\
I &\sqsubseteq \forall r^{m+1}. (\text{pos}_B \geq m \rightarrow \forall r. (M_h \rightarrow \square)) \\
I &\sqsubseteq \forall r. \neg Q_i, \quad \text{for all } i \text{ with } 1 \leq i \leq m,
\end{aligned}$$

where $\text{pos}_B = 0$, $\text{pos}_B = i$ and $\text{pos}_B \geq m$ are abbreviations for Boolean combinations of the concept names B_i which make sure that the position is as required.

This completes the construction of \mathcal{O} . The models of \mathcal{O} and the ABox $\{I(a)\}$ are now almost accepting runs of \mathcal{M} on w . What is missing is the following property:

(*) If two cells t and t' of succeeding configurations have the same position, then the h-memory of t must have the same label $Z_{a, q}$ as the p-memory of t' .

We next convert the constructed ontology \mathcal{O} to a collection of labeled ABox-CQ examples, and then add certain further ABox-CQ examples that enforce (*).

Claim 1. There is a collection E' of ABox-CQ examples of size polynomial in $|\mathcal{O}|$ such that

1. every \mathcal{ALC} -forest model of \mathcal{O} and $\{I(a)\}$ can be extended to a model that satisfies Conditions (a) and (b) of Theorem 7 for E' ;
2. every model that satisfies Conditions (a) and (b) of Theorem 7 for E' is a model of \mathcal{O} and $\{I(a)\}$.

Claim 1 in fact can be shown by using exactly the same examples as in the proof of Theorem 9.

Now we extend E' with additional ABox-CQ examples that enforce $(*)$ to obtain the final collection of examples E . To get a first intuition for the following construction, we ask the reader to recall the example of Theorem 9 in Figure 1 together with the fact that any strong homomorphism from q^* to $\mathcal{I}_{\mathcal{A}^*, h, \mathcal{L}}$ maps x to c or d for every model \mathcal{I} . Intuitively, this gadget uses the idea that we may force existential variables to act like answer variables in the sense that they can only bind to the individuals in \mathcal{ALC} -forest models, but not to elements in the trees of those models. It is not hard to see that all variables in a CQ q that have a directed path to an individual in q are forced to act like such an answer variable. Note that this statement is only true for \mathcal{ALC} , where the edges in the trees of a forest model must be directed away from the ABox.

Building upon these fundamental ideas, we now construct the examples that ensure the contraposition of $(*)$. We use an ABox that matches whenever the $Z_{a,q}$ labels of two cells of consecutive configurations do not agree and then use the corresponding query to ensure that their respective positions do not coincide. The ABox consists of multiple components, only connected through a few designated individuals and each of these components matches into a different bit of the positional memory. The components and the query are carefully crafted in such way that all variables act like answer variables and must collectively map into exactly one of these components. Then, using the concepts A and \bar{A} , this choice presents a certificate for the disagreement of both positions.

For all concept names $Z_{a,q}, Z_{a',q'}$ with $Z_{a,q} \neq Z_{a',q'}$, we define the following positive example (\mathcal{A}^*, q^*) with $\mathcal{A}^* = \mathcal{A}_1 \cup \dots \cup \mathcal{A}_m$ and

$$\mathcal{A}_i = \{ \begin{aligned} & r(c_0^i, c_1^i), \dots, r(c_{m-1}^i, c_m^i), r(c_m^i, c), \\ & r(\hat{c}_0^i, \hat{c}_1^i), \dots, r(\hat{c}_{m-1}^i, \hat{c}_m^i), r(\hat{c}_m^i, \hat{c}), \\ & r(c_0^i, a_0^i), r(a_0^i, a_1^i), \dots, r(a_{m-1}^i, a_m^i), r(a_m^i, c^i), \\ & r(\hat{c}_0^i, b_0^i), r(b_0^i, b_1^i), \dots, r(b_{m-1}^i, b_m^i), r(b_m^i, c^i), \\ & r(\hat{c}_0^i, \hat{a}_0^i), r(\hat{a}_0^i, \hat{a}_1^i), \dots, r(\hat{a}_{m-1}^i, \hat{a}_m^i), r(\hat{a}_m^i, \hat{c}^i), \\ & r(c_0^i, \hat{b}_0^i), r(\hat{b}_0^i, \hat{b}_1^i), \dots, r(\hat{b}_{m-1}^i, \hat{b}_m^i), r(\hat{b}_m^i, \hat{c}^i), \\ & A_i(c^i), r(c, d_1^i), r(d_1^i, d_2^i), r(c^i, d_2^i), \\ & A_i(\hat{c}^i), r(\hat{c}, \hat{d}_1^i), r(\hat{d}_1^i, \hat{d}_2^i), r(\hat{c}^i, \hat{d}_2^i), \\ & R(c_0^i), R(\hat{c}_1^i), \\ & r(\hat{c}, \hat{p}), Z_{a',q'}(\hat{p}), M_p(\hat{p}), \\ & Z_{a,q}(c), M_h(c) \end{aligned} \quad (\text{see Figure 3})$$

$$q^* = \{ \begin{aligned} & r(z_0, z_1), \dots, r(z_{m-1}, z_m), r(z_m, c), \\ & r(\hat{z}_0, \hat{z}_1), \dots, r(\hat{z}_{m-1}, \hat{z}_m), r(\hat{z}_m, \hat{c}), \\ & r(z_0, x_0), r(x_0, x_1), \dots, r(x_{m-1}, x_m), r(x_m, z), \\ & r(\hat{z}_0, \hat{y}_0), r(\hat{y}_0, \hat{y}_1), \dots, r(\hat{y}_{m-1}, \hat{y}_m), r(\hat{y}_m, z), \\ & r(\hat{z}_0, \hat{x}_0), r(\hat{x}_0, \hat{x}_1), \dots, r(\hat{x}_{m-1}, \hat{x}_m), r(\hat{x}_m, \hat{z}), \\ & r(z_0, y_0), r(y_0, y_1), \dots, r(y_{m-1}, y_m), r(y_m, \hat{z}), \\ & A(z), \bar{A}(\hat{z}) \end{aligned} \quad (\text{see Figure 4})$$




Figure 3: The ABox $\mathcal{A}_i \subseteq \mathcal{A}^*$, where arrows depict r -roles.




Figure 4: The query q^* , where arrows depict r -roles.

Observe that the two ABoxes $\mathcal{A}_i, \mathcal{A}_j$ with $i \neq j$ have only three individuals in common, namely c, \hat{c} and \hat{p} . Adding each of these positive examples to E' , we obtain E .

We are now able to state the intended behaviour of (\mathcal{A}^*, q^*) in a precise way.

Claim 2. Let \mathcal{I} be an \mathcal{ALC} -forest model of \mathcal{A}^* . If h is a (strong) homomorphism from q^* to \mathcal{I} , then

1. $h(x) \in \text{ind}(\mathcal{A}^*)$ for all $x \in \text{var}(q^*)$,
2. $\{h(z), h(\hat{z})\} \subseteq \{c^i, \hat{c}^i\}$ for some i with $1 \leq i \leq m$.

Proof of Claim 2. We show both points at the same time. By definition of strong homomorphisms, $h(c) = c \in \text{ind}(\mathcal{A}^*)$ and $h(\hat{c}) = \hat{c} \in \text{ind}(\mathcal{A}^*)$. As \mathcal{ALC} -forest models do not contain predecessors of individuals in \mathcal{A}^* that do not occur in \mathcal{A}^* , $h(z_k) = c_k^i$ and $h(\hat{z}_k) = \hat{c}_k^j$ for all k with $0 \leq k \leq m$ and for some i, j with $1 \leq i, j \leq m$. Observe now that q^* contains a directed r -path of length $m+2$ from z_0 to z and a directed r -path of length $m+2$ from \hat{z}_0 to \hat{z} . Again due to the structure of \mathcal{ALC} -forest models of \mathcal{A}^* , c_0^i and \hat{c}_0^j only have common r -successors at depth $m+2$ for $i = j$, and these are c^i and \hat{c}^i . Thus, $i = j$, and $h(x_k) = a_k^i$ and $h(\hat{y}_k) = b_k^i$, or $h(x_k) = \hat{b}_k^i$ and $h(\hat{y}_k) = \hat{a}_k^i$, for all k with $0 \leq k \leq m$, and therefore $h(z) \in \{c^i, \hat{c}^i\}$. Thus the same argument applies to \hat{z} and the variables y_k and \hat{x}_k . This completes the proof of the claim.

Now consider the intended \mathcal{ALC} -forest model \mathcal{I} of \mathcal{O} depicted in Figure 2, and a homomorphism h from \mathcal{A}^* to \mathcal{I} . By the tree-shape of \mathcal{I} , $h(c_0^i) = h(\hat{c}_0^i)$ and $h(c_0^i) = h(\hat{c}_0^j)$ for all i, j with $1 \leq i, j \leq m$, and $h(c_0^i)$ must be the root (labeled with R) of a configuration K . From the lengths of r -paths, it follows that $h(c)$ is a h-memory node of some cell in this configuration, and from the paths involving d_1^i and d_2^i , it follows that $h(c) = h(a_m^i) = h(b_m^i)$, and $h(c^i)$ is the element that stores the i -th bit of the cell position.

Again the lengths of r -paths imply that $h(\hat{p})$ is the p-memory node of some cell in a successor configuration K' of K . And by the same argument above it follows that $h(\hat{c}^i)$ is the element that stores the i -th bit of this cell position.

Together with Claim 2, these observations suffice to show that this reduction is correct:

1. \mathcal{M} accepts w
2. There exists an \mathcal{ALC} -ontology that fits E

Proof. “1 \Rightarrow 2” Assume there is an accepting computation of \mathcal{M} on w , and let \mathcal{I}' be the corresponding model of \mathcal{O} and $\{I(a)\}$. By construction of the examples E' , \mathcal{I}' can then, as in the proof of Theorem 9, be extended to an interpretation \mathcal{I} that satisfies Conditions (a) and (b) of Theorem 7 for the examples E' . Now we verify that \mathcal{I} also satisfies Condition (b) for the positive examples we added to construct E . Let (\mathcal{A}^*, q^*) be a positive example as constructed above for some $Z_{a,q}, Z_{a',q'}$ with $Z_{a,q} \neq Z_{a',q'}$, and let h be a homomorphism from \mathcal{A}^* to \mathcal{I} . By the above observations about \mathcal{A}^* , it follows that $h(c)$ is a cell labeled with $Z_{a,q}$, and for all i , $h(c^i)$ is the element that stores the i -th bit of the position of this cell. Furthermore, $h(\hat{c})$ is a cell in a successor configuration, where the p-memory is labeled with $Z_{a',q'}$ and for all

i , $h(\hat{c}^i)$ is the element that stores the i -th bit of the position of this cell.

As \mathcal{I} represents a well-formed computation tree, the p-memory of a cell is always labeled with the same $Z_{a,q}$ as the h-memory of the same cell in the predecessor configuration as in (*), it follows that $h(c)$ and $h(\hat{c})$ must be in cells with different cell positions. Thus, there is an i such that $h(c^i) \in A^{\mathcal{I}}$ and $h(\hat{c}^i) \in \bar{A}^{\mathcal{I}}$, or vice versa. This in turn implies $c^i \in A^{\mathcal{I}_{\mathcal{A}^*, h, \mathcal{ALC}}}$ and $\hat{c}^i \in \bar{A}^{\mathcal{I}_{\mathcal{A}^*, h, \mathcal{ALC}}}$, or vice versa. Thus, we can directly construct a (strong) homomorphism g from q^* to $\mathcal{I}_{\mathcal{A}^*, h, \mathcal{ALC}}$ with $g(z) = c^i$ and $g(\hat{z}) = \hat{c}^i$, or vice versa, witnessing $\mathcal{I}_{\mathcal{A}^*, h, \mathcal{ALC}} \models q^*$, as required.

“2 \Rightarrow 1” By Theorem 7 there is an \mathcal{ALC} -forest model \mathcal{I} that satisfies Conditions (a) and (b) for the examples E . Since \mathcal{I} also satisfies Conditions (a) and (b) for the examples E' , \mathcal{I} is a model of \mathcal{O} and $\{I(a)\}$. Thus, \mathcal{I} represents an accepting computation tree of \mathcal{M} on w if it satisfies (*). We will show the contraposition of (*). Thus, let t and t' be cells of succeeding configurations. If the h-memory of t is labeled with $Z_{a,q}$ and the p-memory of t' is labeled with $Z_{a',q'}$ and $Z_{a,q} \neq Z_{a',q'}$, then there is a positive example (\mathcal{A}^*, q^*) constructed for $Z_{a,q}$ and $Z_{a',q'}$ in E and there is a homomorphism h from \mathcal{A}^* to \mathcal{I} with $h(c)$ lies in t and $h(\hat{c})$ lies in t' . As \mathcal{I} satisfies Condition (b) for the example (\mathcal{A}^*, q^*) , it follows that $\mathcal{I}_{\mathcal{A}^*, h, \mathcal{ALC}} \models q^*$. Thus, there is a (strong) homomorphism g from q^* to $\mathcal{I}_{\mathcal{A}^*, h, \mathcal{ALC}}$. As $\mathcal{I}_{\mathcal{A}^*, h, \mathcal{ALC}}$ is an \mathcal{ALC} -forest model of \mathcal{A}^* , $g(z), g(\hat{z}) \subseteq \{c^i, \hat{c}^i\}$ for some i by Claim 2. Since $\bar{A} \sqsubseteq \neg A \in \mathcal{O}$, and \mathcal{I} is a model of \mathcal{O} , $g(z) = c^i$ and $g(\hat{z}) = \hat{c}^i$ or vice versa. Thus $c^i \in A^{\mathcal{I}_{\mathcal{A}^*, h, \mathcal{ALC}}}$ and $\hat{c}^i \in \bar{A}^{\mathcal{I}_{\mathcal{A}^*, h, \mathcal{ALC}}}$, or vice versa. In both cases, the position of t and the position of t' must differ at bit i . Hence, \mathcal{I} satisfies (*) and represents an accepting computation tree of \mathcal{M} on w . \square

G Proofs for Section 7

Let \mathcal{I}_1 and \mathcal{I}_2 be interpretations. We write $(\mathcal{I}_1, d_1) \sim_{\mathcal{ALCQ}} (\mathcal{I}_2, d_2)$ for $d_1 \in \Delta^{\mathcal{I}_1}$ and $d_2 \in \Delta^{\mathcal{I}_2}$ if there exists a *counting bisimulation* $R \subseteq \Delta^{\mathcal{I}_1} \times \Delta^{\mathcal{I}_2}$ with $(d_1, d_2) \in R$. For formal definitions and the following proposition, we refer to (Lutz, Piro, and Wolter 2011).

Proposition 8. *Let \mathcal{I}_1 and \mathcal{I}_2 be interpretations, and \mathcal{O} an \mathcal{ALCQ} -ontology. If $\mathcal{I}_2 \models \mathcal{O}$ and for each $d_1 \in \Delta^{\mathcal{I}_1}$ there exists $d_2 \in \Delta^{\mathcal{I}_2}$ such that $(\mathcal{I}_1, d_1) \sim_{\mathcal{ALCQ}} (\mathcal{I}_2, d_2)$, then \mathcal{I}_1 is a model of \mathcal{O} .*

In the context of \mathcal{ALCQ} , the *unraveling* \mathcal{I}_d of an interpretation \mathcal{I} at $d \in \Delta^{\mathcal{I}}$ is defined analogously to the \mathcal{ALC} case. It is well known that for each $d' \in \Delta^{\mathcal{I}_d}$ there exists $e \in \Delta^{\mathcal{I}}$ such that $(\mathcal{I}_d, d') \sim_{\mathcal{ALCQ}} (\mathcal{I}, e)$.

We say that a homomorphism h from an ABox \mathcal{A}_1 to an ABox \mathcal{A}_2 is *locally injective* if $h(b) \neq h(c)$ for all $r(a, b), r(a, c) \in \mathcal{A}_1$ (Funk et al. 2019).

Proposition 9. *Let \mathcal{I} be an ABox, \mathcal{O} an \mathcal{ALCQ} -ontology and \mathcal{I} an interpretation with $\mathcal{I} \models \mathcal{O}$. If there is a locally injective homomorphism from \mathcal{A} to \mathcal{I} , then there exists a model \mathcal{J} of $\mathcal{O} \cup \mathcal{A}$.*

Proof. Let h be a locally injective homomorphism from \mathcal{A} to \mathcal{I} . \mathcal{J} is then constructed the following way:

We begin with adding $\text{ind}(\mathcal{A})$ to $\Delta^{\mathcal{J}}$ and, for all $a \in \text{ind}(\mathcal{A})$ and $A \in \mathbb{N}_{\mathbb{C}}$, adding a to $A^{\mathcal{J}}$ if $h(a) \in A^{\mathcal{I}}$. Furthermore, for each $a \in \text{ind}(\mathcal{A})$, $r \in \mathbb{N}_{\mathbb{R}}$ and $d \in \Delta^{\mathcal{I}} \setminus \{h(b) \mid r(a, b) \in \mathcal{A}\}$ with $(h(a), d) \in r^{\mathcal{I}}$, add a copy of \mathcal{I}_d to \mathcal{J} and connect a to its root via r . It is easy to see that combining the canonical counting bisimulations of the copies of \mathcal{I}_d with h results in a valid counting bisimulation between \mathcal{J} and \mathcal{I} , and thus for each $e \in \Delta^{\mathcal{J}}$ there exists $e' \in \Delta^{\mathcal{I}}$ such that $(\mathcal{J}, e) \sim_{\mathcal{ALCQ}} (\mathcal{I}, e')$. An application Proposition 8 shows that \mathcal{J} is a model of \mathcal{O} . \square

Theorem 11. Let $E = (E^+, E^-)$ be a collection of labeled ABox examples and $\mathcal{A}^+ = \uplus E^+$. Then the following are equivalent:

1. E admits a fitting \mathcal{ALCQ} -ontology;
2. there is no homomorphism from any $\mathcal{A} \in E^-$ to \mathcal{A}^+ that is locally injective.

Proof. “1 \Rightarrow 2” Assume there exists $\mathcal{A} \in E^-$ and a locally injective homomorphism h from \mathcal{A} to \mathcal{A}^+ . Now suppose, contrary to what we want to show, that there is a fitting \mathcal{ALCQ} -ontology \mathcal{O} . Since \mathcal{O} is consistent with every $\mathcal{A} \in E^+$ it is consistent with \mathcal{A}^+ , witnessed by the disjoint union \mathcal{I} of the individual models. An application of Proposition 9 on \mathcal{A} , \mathcal{I} and the composition $g \circ h$, where g is the homomorphism from \mathcal{A}^+ to \mathcal{I} , proves that \mathcal{A} is consistent with \mathcal{O} . This contradicts the assumption that \mathcal{O} is a fitting ontology for E .

“2 \Rightarrow 1” Assume there is no locally injective homomorphism from any $\mathcal{A} \in E^-$ to \mathcal{A}^+ . We use \mathcal{A}^+ to construct the fitting ontology

$$\mathcal{O} := \mathcal{O}_{\mathcal{A}^+, \Sigma} \cup \bigcup_{a \in \text{ind}(\mathcal{A}^+)} \{\top \sqsubseteq (\leq 1.r)V_a \mid r \in \Sigma\},$$

where $\mathcal{O}_{\mathcal{A}^+, \Sigma}$ is the \mathcal{ALC} -ontology for \mathcal{A}^+ defined in Proposition 6, for $\Sigma = \text{sig}(E)$.

It is straightforward to show that each $\mathcal{A} \in E^+$ is consistent with \mathcal{O} by utilizing $\mathcal{I}_{\mathcal{A}}$, where we set $V_a^{\mathcal{I}_{\mathcal{A}}} := \{a\}$ for each $\text{ind}(\mathcal{A})$. Thus, it remains to prove the inconsistency with every $\mathcal{A} \in E^-$. Contrary to what we want to show, assume that some $\mathcal{A} \in E^-$ is consistent with \mathcal{O} and let \mathcal{I} be a model of \mathcal{A} and \mathcal{O} . Combining \mathcal{I} with Proposition 6, we receive a homomorphism h from \mathcal{A} to \mathcal{A}^+ , where $a \in V_b^{\mathcal{I}}$ if $h(a) = a_b$.

We now claim that h is locally injective. Suppose $r(a, b), r(a, c) \in \mathcal{A}$, then there exists no V_d with $b, c \in V_d^{\mathcal{I}}$, by definition of \mathcal{O} . Lifting this observation back to h shows $h(b) \neq h(c)$, and thus the local injectivity. As the existence of such a homomorphism contradicts the initial assumption, we conclude that \mathcal{O} is a fitting ontology for E . \square