

Conservative Extensions in Horn Description Logics with Inverse Roles

Jean Christoph Jung

JEANJUNG@UNI-BREMEN.DE

Carsten Lutz

CLU@UNI-BREMEN.DE

Mauricio Martel

MAURICIO.MARTEL@GMAIL.COM

Thomas Schneider

THOMAS.SCHNEIDER@UNI-BREMEN.DE

Fachbereich 3 Mathematik/Informatik

Universität Bremen

Postfach 330 440, 28334 Bremen, Germany

Abstract

We investigate the decidability and computational complexity of conservative extensions and the related notions of inseparability and entailment in Horn description logics (DLs) with inverse roles. We consider both query conservative extensions, defined by requiring that the answers to all conjunctive queries are left unchanged, and deductive conservative extensions, which require that the entailed concept inclusions, role inclusions, and functionality assertions do not change. Upper bounds for query conservative extensions are particularly challenging because characterizations in terms of unbounded homomorphisms between universal models, which are the foundation of the standard approach to establishing decidability, fail in the presence of inverse roles. We resort to a characterization that carefully mixes unbounded and bounded homomorphisms and enables a decision procedure that combines tree automata and a mosaic technique. Our main results are that query conservative extensions are 2EXPTIME-complete in all DLs between \mathcal{ELI} and Horn- \mathcal{ALCHIF} and between Horn- \mathcal{ALC} and Horn- \mathcal{ALCHIF} , and that deductive conservative extensions are 2EXPTIME-complete in all DLs between \mathcal{ELI} and \mathcal{ELHIF}_\perp . The same results hold for inseparability and entailment.

1. Introduction

When accessing incomplete data, it can be beneficial to add an ontology formulated in a description logic (DL) to specify relevant domain knowledge, to assign a semantics to the data, and to enrich and unify the vocabulary available for querying. The resulting framework is known as ontology-mediated querying (Bienvenu, ten Cate, Lutz, & Wolter, 2014; Bienvenu & Ortiz, 2015) and, in a data integration context, as ontology-based data access, OBDA (Poggi, Lembo, Calvanese, De Giacomo, Lenzerini, & Rosati, 2008). Significant research activity on ontology-mediated query evaluation has resulted in a thorough understanding of computational complexity trade-offs and in various tools for evaluating queries in practice, for a wide range of DLs such as DL-Lite (Calvanese, De Giacomo, Lembo, Lenzerini, & Rosati, 2007; Artale, Calvanese, Kontchakov, & Zakharyashev, 2009; Rodriguez-Muro & Calvanese, 2012), expressive Horn DLs such as \mathcal{ELI} and Horn- \mathcal{SHIQ} (Hustadt, Motik, & Sattler, 2007; Eiter, Gottlob, Ortiz, & Šimkus, 2008; Eiter, Ortiz, Šimkus, Tran, & Xiao, 2012; Trivela, Stoilos, Chortaras, & Stamou, 2015), and expressive “full Boolean” DLs such as \mathcal{ALC} and \mathcal{SHIQ} (Glimm, Lutz, Horrocks, & Sattler, 2008; Lutz, 2008; Zhou, Cuenca Grau, Nenov, Kaminski, & Horrocks, 2015; Ngo, Ortiz, & Šimkus, 2016).

While query evaluation is by now rather well-understood, there is a need to develop reasoning services that aim at engineering ontologies for ontology-mediated querying and support tasks such as ontology versioning, ontology import, and module extraction, which means to extract a subset from an ontology that is sufficient for the application at hand (Konev, Ludwig, Walther, & Wolter, 2012; Konev, Lutz, Walther, & Wolter, 2009; Grau, Horrocks, Kazakov, & Sattler, 2008). In all these applications, it is important to relate different ontologies. In versioning, for example, one would like to know whether replacing an ontology with a new version has an effect on evaluating the relevant queries. In ontology import, one wants to control the effect on query evaluation of importing an existing ontology. And in module extraction, one wants to know whether the module is really sufficient to evaluate the queries from the application. All this can be formalized by requiring that when exchanging an existing ontology with a new one, the answers to the relevant queries do not change, over all possible data sets (Kontchakov, Wolter, & Zakharyashev, 2010). One arrives at notions of “equivalence” between ontologies that are different from logical equivalence.

We now make this more precise. In DLs, ontologies are represented as a *TBox* while data sets are stored in an *ABox*. A *signature* is a set of concept names and role names. We say that a TBox $\mathcal{T}_2 \supseteq \mathcal{T}_1$ is an (\mathbf{A}, \mathbf{Q}) -*query conservative extension* of a TBox \mathcal{T}_1 , where \mathbf{A} and \mathbf{Q} are signatures relevant for the data and queries, respectively, if all \mathbf{Q} -queries give the same answers w.r.t. \mathcal{T}_1 and w.r.t. \mathcal{T}_2 , for every \mathbf{A} -ABox. We thus identify the relevant queries by signature. Note that the subset relationship $\mathcal{T}_2 \supseteq \mathcal{T}_1$ is natural in some applications such as ontology import and module extraction. It is not natural in other applications such as versioning. In the general case, when \mathcal{T}_1 need not be a subset of \mathcal{T}_2 and the above condition is satisfied, we call \mathcal{T}_1 and \mathcal{T}_2 (\mathbf{A}, \mathbf{Q}) -*query inseparable*. We also consider the notion of query entailment: \mathcal{T}_1 (\mathbf{A}, \mathbf{Q}) -*query entails* \mathcal{T}_2 if all \mathbf{Q} -queries give *at least* the answers w.r.t. \mathcal{T}_1 that they give w.r.t. \mathcal{T}_2 , over every \mathbf{A} -ABox. Clearly, query inseparability and conservative extensions are special cases of query entailment: inseparability is bidirectional entailment and conservative extensions are entailment under the promise that $\mathcal{T}_1 \subseteq \mathcal{T}_2$. When studying decidability or computational complexity, it thus suffices to prove upper bounds for query entailment and lower bounds for conservative extensions. For the query language, we concentrate on conjunctive queries (CQs); since we work with Horn DLs and quantify over the queries, this is equivalent to using unions of CQs (UCQs) and positive existential queries (PEQs). Conservative extensions, inseparability, query entailment of *TBoxes*, as defined above, are useful when knowledge is considered static and data changes frequently. Variants of these notions for *knowledge bases (KBs)*, which consist of a TBox and an ABox, can be used for applications with static data (Wang, Wang, Topor, Pan, & Antoniou, 2014; Arenas, Botoeva, Calvanese, & Ryzhikov, 2016).

CQ entailment has been studied for various DLs (Kontchakov, Pulina, Sattler, Schneider, Selmer, Wolter, & Zakharyashev, 2009; Lutz & Wolter, 2010; Konev et al., 2012; Botoeva, Lutz, Ryzhikov, Wolter, & Zakharyashev, 2019), also in the KB version (Botoeva, Kontchakov, Ryzhikov, Wolter, & Zakharyashev, 2016b; Botoeva et al., 2019), and also for OBDA specifications that involve mappings between data sources and the ABox (Bienvenu & Rosati, 2015), see also the survey by Botoeva, Konev, Lutz, Ryzhikov, Wolter, and Zakharyashev (2016a). Nevertheless, there is still a notable gap in our understanding of this notion: query entailment between TBoxes is poorly understood in Horn DLs with inverse roles, which are considered a crucial feature in many applications. There is in fact a reason

for this: as has already been observed by Botoeva et al. (2016a, 2016b), standard techniques for Horn DLs without inverse roles fail when inverse roles are added. More precisely, for Horn DLs without inverse roles query entailment can be characterized by the existence of homomorphisms between universal models (Lutz & Wolter, 2010; Botoeva et al., 2016a). The resulting characterizations provide an important foundation for decision procedures, often based on tree automata emptiness (Botoeva et al., 2016a). In the presence of inverse roles, however, such characterizations are only correct if one resorts to *bounded* homomorphisms, that is, if one requires the existence of an n -bounded homomorphism, *for any* n (Botoeva et al., 2016a, 2016b). The unbounded n in n -bounded homomorphisms corresponds to CQs of unbounded size that can be used for separating the two TBoxes. Unbounded homomorphisms, in contrast, correspond to infinitary CQs, and while the (implicit) transition to infinitary CQs poses no problems in DLs that do not admit inverse roles, it compromises correctness in the presence of inverse roles. The ‘tighter’ characterization in terms of bounded homomorphism is problematic because it is not obvious how the existence of such (infinite families of) homomorphisms can be verified using tree automata or related techniques and, consequently, decidability results for query conservative extensions in Horn DLs with inverse roles are difficult to obtain. In fact, the only result of which we are aware concerns inseparability of KBs, and it is proved using intricate game-theoretic techniques (Botoeva et al., 2016b).

The aim of this article is to develop decision procedures for and study the complexity of query conservative extensions, query inseparability, and query entailment in Horn DLs with inverse roles such as \mathcal{ELI} and Horn- \mathcal{ALCHIF} . The main idea for establishing decidability is to provide a very careful characterization that mixes unbounded and bounded homomorphisms, pushing the use of bounded homomorphisms to only those places where they cannot possibly be avoided. We can then deal with the part of the characterization that uses unbounded homomorphisms using tree automata while the part that uses bounded homomorphisms is addressed up-front by precomputing relevant information using a mosaic technique. In this way, we establish decidability and a 2EXPTIME upper bound for query entailment in the expressive Horn DL Horn- \mathcal{ALCHIF} , and thus also for query inseparability and query conservative extensions. Together with known lower bounds (Botoeva et al., 2019), this yields 2EXPTIME-completeness for all DLs between \mathcal{ELI} and Horn- \mathcal{ALCHIF} as well as between Horn- \mathcal{ALC} and Horn- \mathcal{ALCHIF} . To be more precise, the complexity is single exponential in the size of \mathcal{T}_1 and double exponential only in the size of \mathcal{T}_2 .

We additionally study the *deductive* version of query entailment, query inseparability, and query conservative extensions. Here, instead of asking whether the answers to all queries are preserved, the question is whether \mathcal{T}_1 entails every concept inclusion, role inclusion, and functionality assertion that \mathcal{T}_2 entails, over a given signature \mathbf{S} . This problem, too, has received considerable interest in the literature, but has not previously been studied for Horn DLs with inverse roles. It is more appropriate than the query-based notions in applications that require conceptual reasoning rather than querying data. It was historically even the first notion of conservative extension studied for DLs (Ghilardi, Lutz, & Wolter, 2006); we again refer to the survey (Botoeva et al., 2016a) for a more detailed discussion. We show that deductive entailment, deductive inseparability, and deductive conservative extensions are 2EXPTIME-complete for all DLs between \mathcal{ELI} and \mathcal{ELHIF}_\perp ; again, the runtime of our algorithm is single exponential in the size of \mathcal{T}_1 and double exponential in the size of \mathcal{T}_2 . For the upper bound, we first show that deductive entailment is essentially identical to query

entailment when the queries are tree-shaped CQs with a single answer variable. We then characterize this version of query entailment using simulations in place of homomorphisms; it is not necessary to resort to a bounded version of simulations. This again enables a decision procedure based on tree automata emptiness. The lower bound is proved by reduction from a certain homomorphism problem between universal models for ABoxes w.r.t. \mathcal{ELI} TBoxes, studied in the context of querying by example (Gutiérrez-Basulto, Jung, & Sabellek, 2018).

This article is structured as follows: in Section 2 we define the notions used throughout the text; in Section 3 we establish the model-theoretic characterizations; in Section 4 we develop the automata-based decision procedure and prove the exact complexity for query entailment, inseparability, and conservative extensions; Section 5 deals with the case of tree-shaped queries and deductive query entailment; Section 6 discusses possible future work. Proofs of some auxiliary lemmas can be found in the appendix.

2. Preliminaries

We define basic notions and lemmas that are needed in the remainder of the article.

2.1 Description Logics

The main DL considered in this article is Horn- \mathcal{ALCHIF} , a member of the Horn- \mathcal{SHIQ} family of DLs whose reasoning problems have been widely studied (Hustadt et al., 2007; Krötzsch, Rudolph, & Hitzler, 2013; Eiter et al., 2008; Kazakov, 2009; Ibáñez-García, Lutz, & Schneider, 2014; Lutz & Wolter, 2017). We introduce Horn- \mathcal{ALCHIF} and several of its fragments. Let $\mathbb{N}_C, \mathbb{N}_R, \mathbb{N}_I$ be countably infinite sets of *concept names*, *role names*, and *individual names*. A *role* is either a role name r or an *inverse role* r^- where r is a role name. As usual, we identify $(r^-)^-$ and r , allowing us to switch between role names and their inverses easily. A *concept inclusion (CI)* is of the form $L \sqsubseteq R$, where L and R are concepts defined by the syntax rules

$$\begin{aligned} R, R' &::= \top \mid \perp \mid A \mid \neg A \mid R \sqcap R' \mid \neg L \sqcup R \mid \exists r.R \mid \forall r.R \\ L, L' &::= \top \mid \perp \mid A \mid L \sqcap L' \mid L \sqcup L' \mid \exists r.L \end{aligned}$$

with A ranging over concept names and r over roles. A *role inclusion (RI)* is of the form $r \sqsubseteq s$ with r, s roles, and a *functionality assertion (FA)* is of the form $\text{func}(r)$ with r a role. A Horn- \mathcal{ALCHIF} TBox \mathcal{T} is a finite set of CIs, RIs, and FAs. To avoid dealing with rather messy technicalities that do neither seem to be very illuminating from a theoretical viewpoint nor too useful from a practical one, we generally assume that functional roles cannot have subroles, that is, $r \sqsubseteq s \in \mathcal{T}$ implies $\text{func}(s) \notin \mathcal{T}$.¹ We conjecture that our main results also hold without that restriction.

In this article, we consider the following fragments of Horn- \mathcal{ALCHIF} . A Horn- \mathcal{ALC} TBox \mathcal{T} is a finite set of CIs that do not use inverse roles. An \mathcal{ELI}_\perp *concept* is an expression that is built according to the syntax rule for L above, but does not use “ \sqcup ”. An \mathcal{ELI}_\perp *CI* is

1. Concerning the usefulness of allowing functional roles to have subroles, we found that only 21 ($\leq 4.8\%$) out of 439 available ontologies in BioPortal (Matentzoglou & Parsia, 2017) contain subroles of functional roles; many of these occurrences appear to be due to modeling mistakes.

a CI of the form $L \sqsubseteq R$ where both L and R are \mathcal{ELI}_\perp concepts. An \mathcal{ELHIF}_\perp TBox is a finite set of \mathcal{ELI}_\perp CIs, RIs, and FAs.

An ABox \mathcal{A} is a finite set of *concept and role assertions* of the form $A(a)$ and $r(a, b)$, where $A \in \mathbf{N}_C$, $r \in \mathbf{N}_R$ and $a, b \in \mathbf{N}_I$. We write $\text{ind}(\mathcal{A})$ for the set of individuals in \mathcal{A} . An ABox \mathcal{A} is *tree-shaped* if (i) \mathcal{A} does not contain an assertion of the form $r(a, a)$, (ii) the undirected graph $G_{\mathcal{A}} = (\text{ind}(\mathcal{A}), \{\{a, b\} \mid r(a, b) \in \mathcal{A}\})$ is a tree, and (iii) there are no multi-edges, that is, for any $a, b \in \text{ind}(\mathcal{A})$, \mathcal{A} contains at most one role assertion that involves both a and b .

The semantics of Horn- \mathcal{ALCHIF} is defined in the usual way (Baader, Horrocks, Lutz, & Sattler, 2017). An *interpretation* is a pair $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, where $\Delta^{\mathcal{I}}$ is a non-empty set, the *domain*, and $\cdot^{\mathcal{I}}$ is the *interpretation function*, with $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$ for every $A \in \mathbf{N}_C$, $r^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$ for every $r \in \mathbf{N}_R$, and $a^{\mathcal{I}} \in \Delta^{\mathcal{I}}$ for every $a \in \mathbf{N}_I$. We make the standard name assumption, that is, $a^{\mathcal{I}} = a$ for all $a \in \mathbf{N}_I$. The interpretation function is extended to inverse roles and arbitrary concepts C in the usual way:

$$\begin{aligned} (r^-)^{\mathcal{I}} &= \{(e, d) \mid (d, e) \in r^{\mathcal{I}}\} \\ \top^{\mathcal{I}} &= \Delta^{\mathcal{I}} \\ \perp^{\mathcal{I}} &= \emptyset \\ (\neg C)^{\mathcal{I}} &= \Delta^{\mathcal{I}} \setminus C^{\mathcal{I}} \\ (C \sqcap D)^{\mathcal{I}} &= C^{\mathcal{I}} \cap D^{\mathcal{I}} \\ (C \sqcup D)^{\mathcal{I}} &= C^{\mathcal{I}} \cup D^{\mathcal{I}} \\ (\exists r.C)^{\mathcal{I}} &= \{d \mid \text{there is } e \in \Delta^{\mathcal{I}} \text{ with } (d, e) \in r^{\mathcal{I}} \text{ and } e \in C^{\mathcal{I}}\} \\ (\forall r.C)^{\mathcal{I}} &= \{d \mid \text{for all } e \in \Delta^{\mathcal{I}}, \text{ if } (d, e) \in r^{\mathcal{I}}, \text{ then } e \in C^{\mathcal{I}}\} . \end{aligned}$$

An interpretation \mathcal{I} satisfies

- a CI $C \sqsubseteq D$, written $\mathcal{I} \models C \sqsubseteq D$, if $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$;
- an RI $r \sqsubseteq s$, written $\mathcal{I} \models r \sqsubseteq s$, if $r^{\mathcal{I}} \subseteq s^{\mathcal{I}}$;
- an FA $\text{func}(r)$, written $\mathcal{I} \models \text{func}(r)$, if $r^{\mathcal{I}}$ is functional;
- a concept assertion $A(a)$, written $\mathcal{I} \models A(a)$, if $a \in A^{\mathcal{I}}$;
- a role assertion $r(a, b)$, written $\mathcal{I} \models r(a, b)$, if $(a, b) \in r^{\mathcal{I}}$.

\mathcal{I} is a *model* of a TBox \mathcal{T} , written $\mathcal{I} \models \mathcal{T}$, if it satisfies all inclusions and assertions in it, likewise for ABoxes. \mathcal{A} is *consistent* with \mathcal{T} if \mathcal{T} and \mathcal{A} have a common model.

Tree-shaped interpretations are defined by analogy with tree-shaped ABoxes. We additionally need a weaker variant that permits multi-edges: an interpretation \mathcal{I} is *weakly tree-shaped* if there are no d, r with $(d, d) \in r^{\mathcal{I}}$ and the undirected graph $(\Delta^{\mathcal{I}}, \{\{d, e\} \mid (d, e) \in r^{\mathcal{I}}\})$ is a tree. Thus, an interpretation \mathcal{I} is *tree-shaped* if it is weakly tree-shaped and has no multi-edges, i.e., $r^{\mathcal{I}} \cap s^{\mathcal{I}} = \emptyset$ for any two distinct roles r, s .

A *signature* \mathbf{S} is a set of concept and role names. An \mathbf{S} -ABox is an ABox that uses only concept and role names from \mathbf{S} , and \mathbf{S} - \mathcal{ELI}_\perp *concepts* and other syntactic objects are defined analogously.

Generally and without further notice, we work with Horn- \mathcal{ALCHIF} TBoxes that are in a certain nesting-free *normal form*, that is, they contain only CIs of the form

$$\top \sqsubseteq A \quad A \sqsubseteq \perp \quad A_1 \sqcap A_2 \sqsubseteq B \quad A \sqsubseteq \exists r.B \quad A \sqsubseteq \forall r.B,$$

where A, B, A_1, A_2 are concept names and r, s are roles. It is well-known that every Horn- \mathcal{ALCHIF} TBox \mathcal{T} can be converted into a TBox \mathcal{T}' in normal form (introducing additional concept names) such that \mathcal{T} is a logical consequence of \mathcal{T}' and every model of \mathcal{T} can be extended to one of \mathcal{T}' by interpreting the additional concept names (see, e.g., Bienvenu, Hansen, Lutz, & Wolter, 2016). As a consequence, all results obtained in this article for TBoxes in normal form lift to the general case. Note that a Horn- \mathcal{ALCHIF} TBox in normal form is essentially an \mathcal{ELI}_\perp TBox since CIs of the form $A \sqsubseteq \forall r.B$ can be equivalently rewritten as $\exists r^-.A \sqsubseteq B$.

2.2 Query Conservative Extensions and Entailment

A *conjunctive query* (CQ) is of the form $q(\mathbf{x}) = \exists \mathbf{y} \varphi(\mathbf{x}, \mathbf{y})$, where \mathbf{x} and \mathbf{y} are tuples of variables and $\varphi(\mathbf{x}, \mathbf{y})$ is a conjunction of *atoms* of the form $A(z)$ or $r(z, z')$ with $A \in \mathbf{N}_C$, $r \in \mathbf{N}_R$, and $z, z' \in \mathbf{x} \cup \mathbf{y}$. We call \mathbf{x} *answer variables* and \mathbf{y} *quantified variables* of q . Tree-shaped and weakly tree-shaped CQs are defined by analogy with (weakly) tree-shaped ABoxes and interpretations: a CQ q is *weakly tree-shaped* if it does not contain atoms of the form $r(z, z)$ and the undirected graph $(\mathbf{x} \cup \mathbf{y}, \{\{z, z'\} \mid r(z, z') \text{ is an atom in } q\})$ is a tree. A weakly tree-shaped CQ q is *tree-shaped* or a *tCQ* if q has no multi-edges, that is, for any two distinct variables z, z' in q , there is at most one atom of the form $r(z, z')$ or $r(z', z)$ in q . A *1tCQ* is a tCQ with exactly one answer variable. We sometimes write $r^-(z, z') \in q$ to mean $r(z', z) \in q$.

A *match* of q in an interpretation \mathcal{I} is a function $\pi : \mathbf{x} \cup \mathbf{y} \rightarrow \Delta^{\mathcal{I}}$ such that $\pi(z) \in A^{\mathcal{I}}$ for every atom $A(z)$ of q and $(\pi(z), \pi(z')) \in r^{\mathcal{I}}$ for every atom $r(z, z')$ of q . We write $\mathcal{I} \models q(a_1, \dots, a_n)$ if there is a match of q in \mathcal{I} with $\pi(x_i) = a_i$ for all $i < n$. A tuple \mathbf{a} of elements from \mathbf{N}_I is a *certain answer* to q over an ABox \mathcal{A} given a TBox \mathcal{T} , written $\mathcal{T}, \mathcal{A} \models q(\mathbf{a})$, if $\mathcal{I} \models q(\mathbf{a})$ for all models of \mathcal{T} and \mathcal{A} .

Definition 1 Let \mathbf{A}, \mathbf{Q} be signatures and $\mathcal{T}_1, \mathcal{T}_2$ Horn- \mathcal{ALCHIF} TBoxes. We say that \mathcal{T}_1 (\mathbf{A}, \mathbf{Q})-CQ entails \mathcal{T}_2 , written $\mathcal{T}_1 \models_{\mathbf{A}, \mathbf{Q}}^{\text{CQ}} \mathcal{T}_2$, if for all \mathbf{A} -ABoxes \mathcal{A} consistent with \mathcal{T}_1 and \mathcal{T}_2 , all \mathbf{Q} -CQs $q(\mathbf{x})$ and all tuples $\mathbf{a} \subseteq \text{ind}(\mathcal{A})$, $\mathcal{T}_2, \mathcal{A} \models q(\mathbf{a})$ implies $\mathcal{T}_1, \mathcal{A} \models q(\mathbf{a})$. If in addition $\mathcal{T}_1 \subseteq \mathcal{T}_2$, we say that \mathcal{T}_2 is an (\mathbf{A}, \mathbf{Q})-CQ conservative extension of \mathcal{T}_1 . If $\mathcal{T}_1 \models_{\mathbf{A}, \mathbf{Q}}^{\text{CQ}} \mathcal{T}_2$ and vice versa, then \mathcal{T}_1 and \mathcal{T}_2 are (\mathbf{A}, \mathbf{Q})-CQ inseparable.

We also consider (\mathbf{A}, \mathbf{Q})-1tCQ entailment, denoted $\models_{\mathbf{A}, \mathbf{Q}}^{1\text{tCQ}}$, (\mathbf{A}, \mathbf{Q})-1tCQ conservative extensions, and (\mathbf{A}, \mathbf{Q})-1tCQ inseparability, defined in the obvious way by replacing CQs with 1tCQs.

If $\mathcal{T}_1 \not\models_{\mathbf{A}, \mathbf{Q}}^{\text{CQ}} \mathcal{T}_2$ because $\mathcal{T}_2, \mathcal{A} \models q(\mathbf{a})$ but $\mathcal{T}_1, \mathcal{A} \not\models q(\mathbf{a})$ for some \mathbf{A} -ABox \mathcal{A} consistent with both \mathcal{T}_i , \mathbf{Q} -CQ $q(\mathbf{x})$ and tuple \mathbf{a} over $\text{ind}(\mathcal{A})$, we call the triple $(\mathcal{A}, q, \mathbf{a})$ a *witness* to non-entailment.

Example 2 Let $\mathcal{T}_1 = \{PhDStud \sqsubseteq \exists advBy.Prof, adv \sqsubseteq advBy^-\}$ and $\mathcal{T}_2 = \mathcal{T}_1 \cup \{\text{func}(advBy)\}$, $\mathbf{A} = \{PhDStud, adv\}$ and $\mathbf{Q} = \{Prof\}$. Then we have $\mathcal{T}_1 \not\models_{\mathbf{A}, \mathbf{Q}}^{\text{CQ}} \mathcal{T}_2$ because of the witness $(\{PhDStud(john), adv(mary, john)\}, Prof(x), mary)$.

If we drop from Definition 1 the condition that \mathcal{A} must be consistent with both \mathcal{T}_1 and \mathcal{T}_2 , then we obtain an alternative notion of CQ entailment that we call *CQ entailment with inconsistent ABoxes*. While CQ entailment with inconsistent ABoxes trivially implies CQ entailment in the original sense, the converse fails.

Example 3 Let $\mathcal{T}_1 = \emptyset$, $\mathcal{T}_2 = \{A_1 \sqcap A_2 \sqsubseteq \perp\}$ and $\mathbf{A} = \{A_1, A_2\}$, $\mathbf{Q} = \{B\}$. Then $\mathcal{T}_1 \models_{\mathbf{A}, \mathbf{Q}}^{\text{CQ}} \mathcal{T}_2$ because none of the \mathcal{T}_i uses the concept name B ; however, \mathcal{T}_1 does not (\mathbf{A}, \mathbf{Q}) -CQ entail \mathcal{T}_2 with inconsistent ABoxes because of the witness $(\{A_1(a), A_2(a)\}, B(x), a)$.

The following lemma relates the two notions of CQ entailment. *CQ evaluation* is the problem to decide, given a TBox \mathcal{T} , an ABox \mathcal{A} , a CQ q , and a tuple $\mathbf{a} \in \text{ind}(\mathcal{A})$, whether $\mathcal{T}, \mathcal{A} \models q(\mathbf{a})$.

Proposition 4 (\mathbf{A}, \mathbf{Q}) -CQ entailment with inconsistent ABoxes can be decided in polynomial time given access to oracles deciding (\mathbf{A}, \mathbf{Q}) -CQ entailment and CQ evaluation.

Consequently and since CQ evaluation is in EXPTIME in Horn- \mathcal{ALCHIF} (Eiter et al., 2008), all complexity results obtained in this article also apply to CQ entailment with inconsistent ABoxes.

To prove Proposition 4, we introduce a notion of entailment that refers to the (in)consistency of ABoxes: \mathcal{T}_1 \mathbf{A} -inconsistency entails \mathcal{T}_2 , written $\mathcal{T}_1 \models_{\mathbf{A}}^{\perp} \mathcal{T}_2$, if for all \mathbf{A} -ABoxes \mathcal{A} : if \mathcal{A} is inconsistent with \mathcal{T}_2 , then \mathcal{A} is inconsistent with \mathcal{T}_1 . Also, given a TBox \mathcal{T} and signatures \mathbf{A}, \mathbf{Q} , we say that \mathcal{T} is (\mathbf{A}, \mathbf{Q}) -universal if $\mathcal{T}, \mathcal{A} \models q(\mathbf{a})$ for all \mathbf{A} -ABoxes \mathcal{A} , \mathbf{Q} -CQs $q(\mathbf{x})$, and tuples \mathbf{a} over $\text{ind}(\mathcal{A})$ that are of the same length as \mathbf{x} . The following is proved in Appendix A.1.

Lemma 5 Let \mathcal{T}_1 and \mathcal{T}_2 be Horn- \mathcal{ALCHIF} TBoxes and let \mathbf{A}, \mathbf{Q} be signatures. Then \mathcal{T}_1 (\mathbf{A}, \mathbf{Q}) -CQ entails \mathcal{T}_2 with inconsistent ABoxes iff one of the two following conditions holds.

- (1) $\mathcal{T}_1 \models_{\mathbf{A}, \mathbf{Q}}^{\text{CQ}} \mathcal{T}_2$ and $\mathcal{T}_1 \models_{\mathbf{A}}^{\perp} \mathcal{T}_2$;
- (2) \mathcal{T}_1 is (\mathbf{A}, \mathbf{Q}) -universal.

Lemma 5 is the core ingredient to proving Proposition 4. We first note that whether \mathcal{T}_1 is (\mathbf{A}, \mathbf{Q}) -universal can be decided by polynomially many CQ evaluation checks. In fact, it is not hard to verify that \mathcal{T}_1 is (\mathbf{A}, \mathbf{Q}) -universal iff \mathbf{Q} contains no role names² and $\mathcal{T}_1, \mathcal{A} \models B(a)$ for all \mathbf{A} -ABoxes \mathcal{A} that contain a single assertion, all \mathbf{Q} -CQs $B(x)$, and all $a \in \text{ind}(\mathcal{A})$. Moreover, $\mathcal{T}_1 \models_{\mathbf{A}}^{\perp} \mathcal{T}_2$ iff the following conditions are satisfied:

1. $\mathcal{T}_1^A \models_{\mathbf{A}, \{A\}}^{\text{CQ}} \mathcal{T}_2^A$, where A is a fresh concept name and \mathcal{T}_i^A is obtained from \mathcal{T}_i by replacing each occurrence of \perp with A and adding the axioms $A \sqsubseteq \forall s.A$ and $A \sqsubseteq \forall s^-.A$ for every role s that occurs in \mathcal{T}_i , for $i = 1, 2$;
2. every \mathbf{A} -ABox $\mathcal{A} = \{r(a, b), r(a, c)\}$ inconsistent with \mathcal{T}_2 is also inconsistent with \mathcal{T}_1 .

This is proved in detail in Appendix A.2. In summary, we obtain Proposition 4.

We close this section with a general remark on the impact of role inclusions on (\mathbf{A}, \mathbf{Q}) -entailment. It is easy to see that $\mathcal{T}_1 \not\models_{\mathbf{A}, \mathbf{Q}}^{\text{CQ}} \mathcal{T}_2$ if there is an \mathbf{A} -role r and a \mathbf{Q} -role s with

2. If there is an $r \in \mathbf{Q}$, then any disconnected ABox \mathcal{A} witnesses that \mathcal{T}_1 is not (\mathbf{A}, \mathbf{Q}) -universal since $\mathcal{T}_1, \mathcal{A} \not\models r(a, b)$ whenever a and b are from different maximal connected components of \mathcal{A} .

$\mathcal{T}_2 \models r \sqsubseteq s$ but $\mathcal{T}_1 \not\models r \sqsubseteq s$. We write $\mathcal{T}_1 \models_{\mathbf{A}, \mathbf{Q}}^{\text{RI}} \mathcal{T}_2$ if there are no such r and s . Clearly, $\mathcal{T}_1 \models_{\mathbf{A}, \mathbf{Q}}^{\text{RI}} \mathcal{T}_2$ can be decided in exponential time: for all $\mathcal{O}(|\mathbf{A}| \cdot |\mathbf{Q}|)$ many pairs of roles, subsumption w.r.t. both \mathcal{T}_i needs to be tested, and each such test can be reduced to concept subsumption (Horrocks & Patel-Schneider, 2004), which is in EXPTIME for the extension *SHIQ* of Horn-*ALCHIF* (Tobies, 2001). It is thus safe to assume $\mathcal{T}_1 \models_{\mathbf{A}, \mathbf{Q}}^{\text{RI}} \mathcal{T}_2$ when deciding CQ entailment, which we will generally do from now on to avoid dealing with special cases.

2.3 Deductive Conservative Extensions

Another natural notion of entailment that has received significant attention in the literature is deductive entailment, which generalizes the notion of deductive conservative extensions, and which separates two TBoxes in terms of concept and role inclusions and functionality assertions, instead of ABoxes and queries (Ghilardi et al., 2006; Lutz, Walther, & Wolter, 2007; Konev et al., 2009; Lutz & Wolter, 2010).

Definition 6 *Let \mathbf{S} be a signature and let \mathcal{T}_1 and \mathcal{T}_2 be \mathcal{ELHIF}_\perp TBoxes. We say that \mathcal{T}_1 \mathbf{S} -deductively entails \mathcal{T}_2 , written $\mathcal{T}_1 \models_{\mathbf{S}}^{\text{TBox}} \mathcal{T}_2$, if for all \mathbf{S} - \mathcal{ELI}_\perp -concept inclusions α and all \mathbf{S} -RIs and \mathbf{S} -FAs $\alpha: \mathcal{T}_2 \models \alpha$ implies $\mathcal{T}_1 \models \alpha$. If additionally $\mathcal{T}_1 \subseteq \mathcal{T}_2$, then we say that \mathcal{T}_2 is an \mathbf{S} -deductive conservative extension of \mathcal{T}_1 . If $\mathcal{T}_1 \models_{\mathbf{S}}^{\text{TBox}} \mathcal{T}_2$ and vice versa, then \mathcal{T}_1 and \mathcal{T}_2 are \mathbf{S} -deductively inseparable.*

Deductive conservative extensions in other DLs is defined accordingly, that is, separation takes place in terms of the TBox statements that are admitted in the DL under consideration. In contrast to (\mathbf{A}, \mathbf{Q}) -query entailment, only one signature is relevant for \mathbf{S} -deductive entailment (unless one would want to distinguish the signatures of the two sides in CIs or RIs, which seems unintuitive). Although \mathbf{S} -deductive entailment and (\mathbf{S}, \mathbf{S}) -query entailment are closely related, it is not difficult to see that they are orthogonal.

Example 7 *Let $\mathcal{T}_1, \mathcal{T}_2$ be as in Example 3 and $\mathbf{S} = \{A_1, A_2, B\}$. Then $\mathcal{T}_1 \not\models_{\mathbf{S}}^{\text{TBox}} \mathcal{T}_2$, witnessed by the \mathcal{ELI} CI $A_1 \sqcap A_2 \sqsubseteq B$. However, $\mathcal{T}_1 \models_{\mathbf{S}, \mathbf{S}}^{\text{CQ}} \mathcal{T}_2$ because \mathbf{S} -CQs cannot detect the disjointness of A_1 and A_2 based on ABoxes that are consistent with \mathcal{T}_2 .*

For the converse direction, let $\mathcal{T}_1 = \emptyset$ and $\mathcal{T}_2 = \{A \sqsubseteq \exists r.B\}$, and $\mathbf{S} = \{A, B\}$. Then $\mathcal{T}_1 \not\models_{\mathbf{S}, \mathbf{S}}^{\text{CQ}} \mathcal{T}_2$ is witnessed by $(\{A(a)\}, \exists x B(x), a)$. However, since \mathcal{ELI}_\perp CIs cannot jump to an unreachable point like the existential quantifier in the CQ $\exists x B(x)$, we have $\mathcal{T}_1 \models_{\mathbf{S}}^{\text{TBox}} \mathcal{T}_2$.

For \mathbf{S} -deductive entailment and (\mathbf{S}, \mathbf{S}) -1tCQ entailment, the connection is even more intimate. In fact, we have the following.

Proposition 8 *In \mathcal{ELHIF}_\perp , \mathbf{S} -deductive entailment can be decided in polynomial time given access to oracles for (\mathbf{S}, \mathbf{S}) -1tCQ entailment and 1tCQ evaluation.*

In fact, Proposition 8 is a consequence of the following lemma and the observations about $\mathcal{T}_1 \models_{\mathbf{S}}^\perp \mathcal{T}_2$ made in Section 2.2. A proof is in Appendix A.3.

Lemma 9 *Let \mathbf{S} be a signature and $\mathcal{T}_1, \mathcal{T}_2 \in \mathcal{ELHIF}_\perp$ TBoxes such that $\mathcal{T}_1 \models_{\mathbf{S}, \mathbf{S}}^{\text{RI}} \mathcal{T}_2$. Then*

$$\mathcal{T}_1 \models_{\mathbf{S}}^{\text{TBox}} \mathcal{T}_2 \quad \text{iff} \quad \mathcal{T}_1 \models_{\mathbf{S}, \mathbf{S}}^{\text{1tCQ}} \mathcal{T}_2 \quad \text{and} \quad \mathcal{T}_1 \models_{\mathbf{S}}^\perp \mathcal{T}_2.$$

Note that Definition 6 assumes the TBoxes \mathcal{T}_1 and \mathcal{T}_2 to be formulated in \mathcal{ELHIF}_\perp rather than in Horn- \mathcal{ALCHIF} . This is because the main theme of this article are query conservative extensions and, as demonstrated by Proposition 8, separation in terms of \mathcal{ELI}_\perp concept inclusions is closely related to separation by (tree-shaped) ABoxes and ItCQs. Now, separation in terms of \mathcal{ELI}_\perp concept inclusions is the natural choice for \mathcal{ELHIF}_\perp TBoxes as \mathcal{ELHIF}_\perp is based on the concept language \mathcal{ELI}_\perp while this is not true for Horn- \mathcal{ALCHIF} . All results in this article that concern deductive conservative extensions (and related notions) actually remain valid also when \mathcal{T}_1 and \mathcal{T}_2 are Horn- \mathcal{ALCHIF} TBoxes, while separation is still in terms of \mathcal{ELI}_\perp CIs. Separating Horn- \mathcal{ALCHIF} TBoxes in terms of Horn- \mathcal{ALCHIF} is an interesting topic that is outside the scope of this article, see Section 6 for some additional discussion.

2.4 Homomorphisms and Simulations

Homomorphisms, both in their (standard) unbounded form and in bounded form, are an elementary tool for dealing with CQs. Likewise, (unbounded) simulations are closely linked to ItCQs. We will make extensive use of these notions throughout the article.

For interpretations $\mathcal{I}_1, \mathcal{I}_2$ and a signature \mathbf{S} , an \mathbf{S} -homomorphism from \mathcal{I}_1 to \mathcal{I}_2 is a total function $h : \Delta^{\mathcal{I}_1} \rightarrow \Delta^{\mathcal{I}_2}$ that satisfies the following conditions.

- (1) $h(a) = a$ for all $a \in \mathbf{N}_I$;
- (2) $h(d) \in A^{\mathcal{I}_2}$ for all $d \in A^{\mathcal{I}_1}$, $A \in \mathbf{N}_C \cap \mathbf{S}$;
- (3) $(h(d), h(d')) \in r^{\mathcal{I}_2}$ for all $(d, d') \in r^{\mathcal{I}_1}$, $r \in \mathbf{N}_R \cap \mathbf{S}$.

We write $\mathcal{I}_1 \rightarrow_{\mathbf{S}} \mathcal{I}_2$ to denote the existence of an \mathbf{S} -homomorphism from \mathcal{I}_1 to \mathcal{I}_2 . If $\mathbf{S} = \mathbf{N}_C \cup \mathbf{N}_R$, we write $\mathcal{I} \rightarrow \mathcal{J}$.

Let $\mathcal{I}_1, \mathcal{I}_2$ be interpretations, $d \in \Delta^{\mathcal{I}_1}$, and $n \geq 0$. We say that there is an n -bounded \mathbf{S} -homomorphism from \mathcal{I}_1 to \mathcal{I}_2 , written $\mathcal{I}_1 \rightarrow_{\mathbf{S}}^n \mathcal{I}_2$, if for any subinterpretation³ \mathcal{I}'_1 of \mathcal{I}_1 with $|\Delta^{\mathcal{I}'_1}| \leq n$, we have $\mathcal{I}'_1 \rightarrow_{\mathbf{S}} \mathcal{I}_2$. Moreover, we write $\mathcal{I}_1 \rightarrow_{\mathbf{S}}^{\text{fin}} \mathcal{I}_2$ if $\mathcal{I}_1 \rightarrow_{\mathbf{S}}^n \mathcal{I}_2$ for every n . If $\mathbf{S} = \mathbf{N}_C \cup \mathbf{N}_R$, we write $\mathcal{I} \rightarrow^{\text{fin}} \mathcal{J}$.

For $d \in \Delta^{\mathcal{I}}$ and $n \geq 0$, we denote with $\mathcal{I}|_n^d$ the restriction of interpretation \mathcal{I} to elements that can be reached by starting at d and traveling along at most n role edges (forwards or backwards). The following is standard to prove.

Lemma 10 *Let \mathbf{S} be a signature and $\mathcal{I}_1, \mathcal{I}_2$ be interpretations.*

1. *If $\mathcal{I}_1 \rightarrow_{\mathbf{S}} \mathcal{I}_2$, then for all \mathbf{S} -CQs q and tuples \mathbf{a} , $\mathcal{I}_1 \models q(\mathbf{a})$ implies $\mathcal{I}_2 \models q(\mathbf{a})$.*
2. *For every $n \geq 0$, if $\mathcal{I}_1 \rightarrow_{\mathbf{S}}^n \mathcal{I}_2$, then for all \mathbf{S} -CQs q with at most n variables and all tuples \mathbf{a} , $\mathcal{I}_1 \models q(\mathbf{a})$ implies $\mathcal{I}_2 \models q(\mathbf{a})$.*
3. *If \mathcal{I}_1 is finitely branching, then the following are equivalent:*

- (a) $\mathcal{I}_1 \rightarrow_{\mathbf{S}}^{\text{fin}} \mathcal{I}_2$;

3. We mean an interpretation \mathcal{I}'_1 such that $\Delta^{\mathcal{I}'_1} \subseteq \Delta^{\mathcal{I}}$, $A^{\mathcal{I}'_1} = A^{\mathcal{I}} \cap \Delta^{\mathcal{I}'_1}$ for all concept names A , and $r^{\mathcal{I}'_1} = r^{\mathcal{I}} \cap (\Delta^{\mathcal{I}'_1} \times \Delta^{\mathcal{I}'_1})$ for all role names r .

(b) $\mathcal{I}_1|_i^d \rightarrow_{\mathbf{S}} \mathcal{I}_2$ for every $d \in \Delta^{\mathcal{I}_1}$ and $i \geq 0$.

Given a signature \mathbf{S} and two interpretations \mathcal{I}, \mathcal{J} , an \mathbf{S} -simulation of \mathcal{I} in \mathcal{J} is a relation $\sigma \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{J}}$ that satisfies the following conditions.

- (1) $(a, a) \in \sigma$ for all $a \in \mathbf{N}_1$;
- (2) if $d \in A^{\mathcal{I}}$ with $A \in \mathbf{S}$ and $(d, e) \in \sigma$, then $e \in A^{\mathcal{J}}$;
- (3) if $(d, d') \in r^{\mathcal{I}}$ with r an \mathbf{S} -role and $(d, e) \in \sigma$, then there is some e' with $(e, e') \in r^{\mathcal{J}}$ and $(d', e') \in \sigma$.

We write $\mathcal{I} \preceq_{\mathbf{S}} \mathcal{J}$ if there is an \mathbf{S} -simulation of \mathcal{I} in \mathcal{J} . Point 1 of Lemma 10 still holds when homomorphisms are replaced with simulations and CQs with 1tCQs.

2.5 The Universal Model

We next introduce universal models, whose existence is a distinguishing feature of Horn logics. More precisely, for every Horn- \mathcal{ALCHIF} TBox \mathcal{T} and ABox \mathcal{A} , there is a model $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ of \mathcal{T} and \mathcal{A} that homomorphically embeds into every model of \mathcal{T} and \mathcal{A} . This last property, together with Point 1 of Lemma 10, ensures that $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ is universal in the sense that the certain answers to any CQ q over \mathcal{A} given \mathcal{T} can be obtained by evaluating q over $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$. We define the universal model by extending the ABox in a forward chaining way akin to the chase in database theory.

Let \mathcal{T} be a Horn- \mathcal{ALCHIF} TBox in normal form and \mathcal{A} an ABox that is consistent with \mathcal{T} . For a set t of concept names, we write $\sqcap t$ as a shorthand for $\sqcap_{A \in t} A$. A \mathcal{T} -type is a set t of concept names that occur in \mathcal{T} such that $\mathcal{T} \models \sqcap t \sqsubseteq A$ implies $A \in t$ for all concept names A . For $a \in \text{ind}(\mathcal{A})$ and with \mathcal{T} understood, let $\text{tp}_{\mathcal{T}}(a)$ denote the \mathcal{T} -type $\{A \mid \mathcal{T}, \mathcal{A} \models A(a)\}$. When t, t' are \mathcal{T} -types, r is a role, and $a \in \text{ind}(\mathcal{A})$, we write

- $t \rightsquigarrow_r^{\mathcal{T}} t'$ if $\mathcal{T} \models \sqcap t \sqsubseteq \exists r. \sqcap t'$ and t' is maximal with this condition;
- $a \rightsquigarrow_r^{\mathcal{T}, \mathcal{A}} t$ if $\mathcal{T}, \mathcal{A} \models \exists r. \sqcap t(a)$ and t is maximal with this condition.

The relations $\rightsquigarrow_r^{\mathcal{T}}$ and $\rightsquigarrow_r^{\mathcal{T}, \mathcal{A}}$ can be computed in exponential time: for every t, t' (resp. a, t), the above entailment test requires at most exponential time since subsumption (resp. instance checking) in \mathcal{SHIQ} is EXPTIME-complete (Tobies, 2001). A *path for \mathcal{A} and \mathcal{T}* is a finite sequence $\pi = ar_0t_1 \cdots t_{n-1}r_{n-1}t_n$, $n \geq 0$, with $a \in \text{ind}(\mathcal{A})$, r_0, \dots, r_{n-1} roles, and t_1, \dots, t_n \mathcal{T} -types such that

- (i) $a \rightsquigarrow_{r_0}^{\mathcal{T}, \mathcal{A}} t_1$ and, if $\text{func}(r_0) \in \mathcal{T}$, then there is no $b \in \text{ind}(\mathcal{A})$ such that $\mathcal{T}, \mathcal{A} \models r_0(a, b)$;
- (ii) $t_i \rightsquigarrow_{r_i}^{\mathcal{T}} t_{i+1}$ and, if $\text{func}(r_i) \in \mathcal{T}$, then $r_{i-1} \neq r_i^-$, for $1 \leq i < n$.

When $n > 0$, we use $\text{tail}(\pi)$ to denote t_n . Let Paths be the set of all paths for \mathcal{A} and \mathcal{T} . The *universal model* $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ of \mathcal{T} and \mathcal{A} is defined as follows:

$$\begin{aligned} \Delta^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}} &= \text{Paths} \\ A^{\mathcal{I}_{\mathcal{T}, \mathcal{A}}} &= \{a \in \text{ind}(\mathcal{A}) \mid \mathcal{T}, \mathcal{A} \models A(a)\} \cup \{\pi \in \Delta^{\mathcal{I}} \setminus \text{ind}(\mathcal{A}) \mid \mathcal{T} \models \sqcap \text{tail}(\pi) \sqsubseteq A\} \end{aligned}$$

$$\begin{aligned}
 r^{\mathcal{I}_{\mathcal{T},\mathcal{A}}} = & \{(a, b) \in \text{ind}(\mathcal{A}) \times \text{ind}(\mathcal{A}) \mid s(a, b) \in \mathcal{A} \text{ and } \mathcal{T} \models s \sqsubseteq r\} \cup \\
 & \{(\pi, \pi st) \mid \pi st \in \text{Paths} \text{ and } \mathcal{T} \models s \sqsubseteq r\} \cup \\
 & \{(\pi st, \pi) \mid \pi st \in \text{Paths} \text{ and } \mathcal{T} \models s^- \sqsubseteq r\}.
 \end{aligned}$$

We also need a universal model $\mathcal{I}_{\mathcal{T},t}$ of a TBox \mathcal{T} and a type t , instead of an ABox. More precisely, we define $\mathcal{I}_{\mathcal{T},t} = \mathcal{I}_{\mathcal{T},\mathcal{A}_t}$ where $\mathcal{A}_t = \{A(a) \mid A \in t\}$ for a fixed $a \in \mathbb{N}_1$.

The following lemma summarizes the main properties of universal models. The proof is standard and omitted (see, e.g., Botoeva et al., 2016b).

Lemma 11 *For every Horn-ALCHIF TBox \mathcal{T} in normal form and ABox \mathcal{A} consistent with \mathcal{T} , the following hold:*

- (1) $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ is a model of \mathcal{T} and \mathcal{A} ;
- (2) $\mathcal{I}_{\mathcal{T},\mathcal{A}} \rightarrow \mathcal{I}$ for all models \mathcal{I} of \mathcal{T} and \mathcal{A} ;
- (3) $\mathcal{T}, \mathcal{A} \models q(\mathbf{a})$ iff $\mathcal{I}_{\mathcal{T},\mathcal{A}} \models q(\mathbf{a})$, for all CQs $q(\mathbf{x})$ and tuples \mathbf{a} of individuals;
- (4) $\mathcal{I}_{\mathcal{T},t} \rightarrow \mathcal{I}_{\mathcal{T},t'}$ for all \mathcal{T} -types t, t' with $t \subseteq t'$.

We shall sometimes refer to subinterpretations of universal models. Given a TBox \mathcal{T} , an ABox \mathcal{A} , and $a \in \text{ind}(\mathcal{A})$, we use $\mathcal{I}_{\mathcal{T},\mathcal{A}}|_a$ to denote the subinterpretation of $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ rooted at a , i.e., the restriction of $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ to all paths π beginning with a . Clearly, $\mathcal{I}_{\mathcal{T},\mathcal{A}}|_a$ is weakly tree-shaped.

3. Model-Theoretic Characterization

We aim to provide a model-theoretic characterization of CQ-entailment that will be the basis for our decision procedure later on. We proceed in two steps. We first show that non-entailment between \mathcal{T}_1 and \mathcal{T}_2 is always witnessed by a tree-shaped ABox and a weakly tree-shaped CQ with at most one answer variable. Second, we provide the desired characterization, using a careful mix of bounded and unbounded homomorphisms. We also establish a version for 1tCQ-entailment that uses (unbounded) simulations.

3.1 Tree-Shaped Witnesses

We show that non-entailment between \mathcal{T}_1 and \mathcal{T}_2 is always witnessed by a tree-shaped ABox and a weakly tree-shaped CQ with at most one answer variable. The idea is to first manipulate the witnessing CQ so that it takes the desired form and to then unravel the ABox.

We start with defining ABox unraveling (see also Lutz & Wolter, 2017). The *unraveling* $U_{\mathcal{A}}^a$ of an ABox \mathcal{A} at an individual $a \in \text{ind}(\mathcal{A})$ is the following (usually infinite) ABox:

- $\text{ind}(U_{\mathcal{A}}^a)$ is the set of sequences $b_0 r_0 b_1 \cdots r_{n-1} b_n$ with $n \geq 0$, where $b_0 = a$, $b_i \in \text{ind}(\mathcal{A})$ for all $0 \leq i \leq n$, $r_i(b_i, b_{i+1}) \in \mathcal{A}$ for all $0 \leq i < n$, and $(b_{i-1}, r_{i-1}^-) \neq (b_{i+1}, r_i)$ (the latter inequality is needed to ensure preservation of functionality);

- the concept assertions in $U_{\mathcal{A}}^a$ are all assertions of the shape $C(\alpha)$ such that $\alpha = b_0 \cdots b_n \in \text{ind}(U_{\mathcal{A}}^a)$ and $C(b_n) \in \mathcal{A}$, and the role assertions in $U_{\mathcal{A}}^a$ are all assertions of the shape $r(b_0 \cdots b_{n-1}, \alpha)$ such that $\alpha = b_0 \cdots b_{n-1} r b_n \in \text{ind}(U_{\mathcal{A}}^a)$.

The following is not hard to prove (see Lutz & Wolter, 2017).

Lemma 12 *Let \mathcal{T} be a Horn-ALCHIF TBox, \mathcal{A} an ABox that is consistent with \mathcal{T} , and $a \in \text{ind}(\mathcal{A})$. Then*

- (a) $U_{\mathcal{A}}^a$ is consistent with \mathcal{T} ;
- (b) $\mathcal{I}_{\mathcal{T}, U_{\mathcal{A}}^a} \rightarrow \mathcal{I}_{\mathcal{T}, \mathcal{A}}$;
- (c) $\mathcal{I}_{\mathcal{T}, U_{\mathcal{A}}^a}|_a$ is isomorphic to $\mathcal{I}_{\mathcal{T}, \mathcal{A}}|_a$.

We now prove the main result of this section.

Proposition 13 *Let \mathcal{T}_1 and \mathcal{T}_2 be Horn-ALCHIF TBoxes with $\mathcal{T}_1 \models_{\mathbf{A}, \mathbf{Q}}^{\text{RI}} \mathcal{T}_2$. If $\mathcal{T}_1 \not\models_{\mathbf{A}, \mathbf{Q}}^{\text{CQ}} \mathcal{T}_2$, then there is a witness $(\mathcal{A}, q, \mathbf{a})$ where \mathcal{A} is tree-shaped and q is weakly tree-shaped and has at most one answer variable.*

Proof. Assume $\mathcal{T}_1 \not\models_{\mathbf{A}, \mathbf{Q}}^{\text{CQ}} \mathcal{T}_2$, i.e., $\mathcal{T}_2, \mathcal{A} \models q(\mathbf{a})$ and $\mathcal{T}_1, \mathcal{A} \not\models q(\mathbf{a})$, for some \mathbf{A} -ABox \mathcal{A} consistent with both \mathcal{T}_i , some \mathbf{Q} -CQ $q(\mathbf{x})$ and some tuple \mathbf{a} . Lemma 11 (3) yields $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}} \models q(\mathbf{a})$ and $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}} \not\models q(\mathbf{a})$. We first show that the following properties of q and \mathbf{a} are without loss of generality.

- (a) Every match π of $q(\mathbf{x})$ in $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$ with $\pi(\mathbf{x}) = \mathbf{a}$ maps every quantified variable into the anonymous part;
- (b) every match π of $q(\mathbf{x})$ in $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$ with $\pi(\mathbf{x}) = \mathbf{a}$ maps the quantified variables to pairwise distinct elements;
- (c) $q(\mathbf{x})$ does not contain atoms of the form $r(x_1, x_2)$ with x_1, x_2 answer variables;
- (d) $q(\mathbf{x})$ is connected.

For (a), take a match π of q in $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$ with $\pi(\mathbf{x}) = \mathbf{a}$. If there is a quantified variable y such that $\pi(y) = b \in \text{ind}(\mathcal{A})$, obtain $q'(\mathbf{x}, y)$ from $q(\mathbf{x})$ by removing the quantification over y , thus making y an answer variable. Clearly, we have $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}} \models q'(\mathbf{a}, b)$ and $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}} \not\models q'(\mathbf{a}, b)$, and thus $\mathcal{T}_2, \mathcal{A} \models q'(\mathbf{a}, b)$ and $\mathcal{T}_1, \mathcal{A} \not\models q'(\mathbf{a}, b)$.

For (b), take a match π of q in $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$ with $\pi(\mathbf{x}) = \mathbf{a}$. If there are quantified variables y, y' such that $y \neq y'$ and $\pi(y) = \pi(y')$, obtain $q'(\mathbf{x})$ from $q(\mathbf{x})$ by replacing all occurrences of y' with y and removing quantification over y' . Clearly, we have $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}} \models q'(\mathbf{a})$ and $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}} \not\models q'(\mathbf{a})$, and thus $\mathcal{T}_2, \mathcal{A} \models q'(\mathbf{a})$ and $\mathcal{T}_1, \mathcal{A} \not\models q'(\mathbf{a})$.

For (c), let $q(\mathbf{x}) = \exists \mathbf{y} (r(x_1, x_2) \wedge \varphi(\mathbf{x}', \mathbf{y}))$ with $x_1, x_2 \in \mathbf{x}$, and let $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}} \models q(\mathbf{a})$ be witnessed by the match π with $\pi(x_i) = a_i$, $i = 1, 2$. Construct the CQ $q'(\mathbf{x}') = \exists \mathbf{y} \varphi(\mathbf{x}', \mathbf{y})$ by dropping the atom $r(x_1, x_2)$ (and thus possibly removing x_1 and/or x_2 from the free variables). It is clear that $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}} \models q'(\mathbf{a}')$ for the corresponding restriction \mathbf{a}' of the tuple \mathbf{a} ; thus it suffices to show that $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}} \not\models q'(\mathbf{a}')$.

From $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}} \models q(\mathbf{a})$ we can conclude that $(a_1, a_2) \in r^{\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}}$. By construction of $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$ there is some \mathbf{A} -role r' with $r'(a_1, a_2) \in \mathcal{A}$ and $\mathcal{T}_2 \models r' \sqsubseteq r$ (which includes the possibility $r' = r$, i.e., $r(a_1, a_2) \in \mathcal{A}$). Due to $\mathcal{T}_1 \models_{\mathbf{A}, \mathbf{Q}}^{\text{RI}} \mathcal{T}_2$, we also have $\mathcal{T}_1 \models r' \sqsubseteq r$ and hence $(a_1, a_2) \in r^{\mathcal{I}_{\mathcal{T}_1, \mathcal{A}}}$. This implies the desired $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}} \not\models q'(\mathbf{a}')$ because, otherwise, any match π of q' in $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$ with $\pi(x_i) = a_i$, $i = 1, 2$, could be extended to a match of q .

For (d), observe that $\mathcal{T}_2, \mathcal{A} \models q(\mathbf{a})$ and $\mathcal{T}_1, \mathcal{A} \not\models q(\mathbf{a})$ implies $\mathcal{T}_2, \mathcal{A} \models q'(\mathbf{a})$ and $\mathcal{T}_1, \mathcal{A} \not\models q'(\mathbf{a})$ for some connected component q' of q .

Thus, as long as q violates any of the above properties, we apply the corresponding modification; it is routine to verify that after finitely many steps q satisfies properties (a) to (d). We verify that then q is weakly tree-shaped and has at most one answer variable.

To show the latter, assume that \mathbf{x} contains more than one answer variable, say $x \neq x'$. By (d), q is connected, and thus, there is a path from x to x' in q , that is, a sequence of atoms $r_1(z_1, z_2), r_2(z_2, z_3), \dots, r_n(z_n, z_{n+1})$ with $z_1 = x, z_{n+1} = x'$ and with roles r_i such that $z_{i+1} \neq z_{i-1}$ for every $1 < i \leq n$. Let π be a match of q in $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$ with $\pi(\mathbf{x}) = \mathbf{a}$ and consider z_2 . Because of property (c), z_2 cannot be an answer variable. By property (a), $\pi(z_2)$ is a successor of $\pi(z_1)$ in $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}|_{\pi(z_1)}$. Since π is a match with $\pi(x') \in \text{ind}(\mathcal{A})$ and since $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}|_{\pi(z_1)}$ is weakly tree-shaped, there has to be another element $z_\ell \neq z_2$ such that $\pi(z_\ell) = \pi(z_2)$, in contradiction with property (b).

For the former, assume that q is not weakly tree-shaped, that is, it contains a cycle $r_1(z_1, z_2), r_2(z_2, z_3), \dots, r_n(z_n, z_{n+1})$ such that $z_1 = z_{n+1}, z_2 \neq z_n$, and $z_{i+1} \neq z_{i-1}$ for every $1 < i \leq n$. Let π be a match of $q(\mathbf{x})$ in $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$ with $\pi(\mathbf{x}) = \mathbf{a}$. Since q has at most one answer variable and satisfies properties (a) and (d), there is some ABox individual $a \in \text{ind}(\mathcal{A})$ such that π maps all variables z_1, \dots, z_{n+1} into $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}|_a$, which is weakly tree-shaped by definition. Since π is a match for the cycle in this part, there have to be $i \neq j$ such that z_i, z_j are quantified variables with $z_i \neq z_j$ and $\pi(z_i) = \pi(z_j)$, in contradiction to property (b).

It remains to transform \mathcal{A} into a tree-shaped ABox. Due to properties (a) and (d), we can fix an ABox individual a such that there is a match π of q in $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}|_a$. Note that, if q has a single answer variable x , then $\pi(x) = a$. Consider the unraveling $U_{\mathcal{A}}^a$ of \mathcal{A} at a . Due to Lemma 12 (a), $U_{\mathcal{A}}^a$ is still consistent with both \mathcal{T}_i . Lemma 12 (c) and the fact that the range of π is contained in the domain of $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}|_a$ imply $\mathcal{T}_2, U_{\mathcal{A}}^a \models q(a)$. Moreover, Lemma 12 (b) and $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}} \not\models q(a)$ imply $\mathcal{T}_1, U_{\mathcal{A}}^a \not\models q(a)$. By compactness, there is a finite subset $\mathcal{B} \subseteq U_{\mathcal{A}}^a$ with $\mathcal{T}_2, \mathcal{B} \models q(a)$ and $\mathcal{T}_1, \mathcal{B} \not\models q(a)$. Clearly, we can also assume that \mathcal{B} is connected. \square

We also observe an analogous result for 1tCQ entailment.

Proposition 14 *Let \mathcal{T}_1 and \mathcal{T}_2 be Horn-ALCHIF TBoxes with $\mathcal{T}_1 \models_{\mathbf{A}, \mathbf{Q}}^{\text{RI}} \mathcal{T}_2$. If $\mathcal{T}_1 \not\models_{\mathbf{A}, \mathbf{Q}}^{\text{1tCQ}} \mathcal{T}_2$, then there is a witness (\mathcal{A}, q, a) where \mathcal{A} is tree-shaped and q is a 1tCQ.*

Proof. Let (\mathcal{A}, a, q) be a witness for $\mathcal{T}_1 \not\models_{\mathbf{A}, \mathbf{Q}}^{\text{1tCQ}} \mathcal{T}_2$, that is, $\mathcal{T}_2, \mathcal{A} \models q(a)$ and $\mathcal{T}_1, \mathcal{A} \not\models q(a)$. As q is already a 1tCQ, it suffices to modify the ABox. Consider the unraveling $U_{\mathcal{A}}^a$ of \mathcal{A} at a . We can argue as in the proof of Proposition 13 that there is a tree-shaped (and finite) ABox $\mathcal{B} \subseteq U_{\mathcal{A}}^a$ such that $\mathcal{T}_2, \mathcal{B} \models q(a)$ and $\mathcal{T}_1, \mathcal{B} \not\models q(a)$. \square

3.2 Characterization for CQ Entailment

We now develop the announced characterization of CQ entailment. Known characterizations of entailment in Horn DLs without inverse roles (Lutz & Wolter, 2010; Botoeva et al., 2019) suggest that a first natural candidate characterization would be: $\mathcal{T}_1 \models_{\mathbf{A}, \mathbf{Q}}^{\text{CQ}} \mathcal{T}_2$ iff $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}} \rightarrow_{\mathbf{Q}} \mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$ for all tree-shaped \mathbf{A} -ABoxes \mathcal{A} that are consistent with \mathcal{T}_1 and \mathcal{T}_2 . However, such a characterization fails in the presence of inverse roles, in fact already for \mathcal{ELI} .

Example 15 *Let $\mathcal{T}_1 = \{A \sqsubseteq \exists s.B, B \sqsubseteq \exists r^-.B\}$, $\mathcal{T}_2 = \{A \sqsubseteq \exists s.B, B \sqsubseteq \exists r.B\}$, $\mathbf{A} = \{A\}$, and $\mathbf{Q} = \{r\}$. Then both $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$ and $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$ contain an infinite r -path; the r -path in $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$ has a final element while the one in $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$ does not. Hence, we have $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}} \not\rightarrow_{\mathbf{Q}} \mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$. However, it can be verified that $\mathcal{T}_1 \models_{\mathbf{A}, \mathbf{Q}}^{\text{CQ}} \mathcal{T}_2$, see Theorem 18 below.*

It is easy to turn the above candidate characterization into a correct one by replacing homomorphisms with bounded homomorphisms, cf. Points 1 and 2 of Lemma 10. Since we are interested in *all* CQs over signature \mathbf{Q} , however, we have to replace unbounded homomorphisms with n -bounded homomorphisms, for *any* bound n . In fact, the following characterization follows from the definition of CQ entailment, Proposition 13, and Lemma 10.

Lemma 16 *Let \mathcal{T}_1 and \mathcal{T}_2 be Horn-ALCHIF TBoxes with $\mathcal{T}_1 \models_{\mathbf{A}, \mathbf{Q}}^{\text{RI}} \mathcal{T}_2$. Then $\mathcal{T}_1 \models_{\mathbf{A}, \mathbf{Q}}^{\text{CQ}} \mathcal{T}_2$ iff for all tree-shaped \mathbf{A} -ABoxes \mathcal{A} consistent with \mathcal{T}_1 and \mathcal{T}_2 , we have $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}} \rightarrow_{\mathbf{Q}}^{\text{fin}} \mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$.*

Proof. We prove both implications via contraposition.

“ \Leftarrow ”. Assume $\mathcal{T}_1 \not\models_{\mathbf{A}, \mathbf{Q}}^{\text{CQ}} \mathcal{T}_2$ and consider a witness $(\mathcal{A}, q, \mathbf{a})$. By Proposition 13, we can assume that \mathcal{A} is tree-shaped. From Lemma 11 (3), we get $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}} \models q(\mathbf{a})$ and $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}} \not\models q(\mathbf{a})$. Let π be a match for $q(\mathbf{x})$ in $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$ with $\pi(\mathbf{x}) = (\mathbf{a})$, and let \mathcal{I} be the finite subinterpretation of $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$ whose domain $\Delta^{\mathcal{I}}$ is the range of π . We have $\mathcal{I} \not\rightarrow_{\mathbf{Q}} \mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$ because of $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}} \not\models q(\mathbf{a})$ and $\mathcal{I} \models q(\mathbf{a})$. Hence $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}} \not\rightarrow_{\mathbf{Q}}^{\text{fin}} \mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$.

“ \Rightarrow ”. Assume $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}} \not\rightarrow_{\mathbf{Q}}^{\text{fin}} \mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$, that is, there is a finite subinterpretation \mathcal{I} of $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$ with $\mathcal{I} \not\rightarrow_{\mathbf{Q}} \mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$. We associate a query $q_{\mathcal{I}}$ with \mathcal{I} as follows. The variables in $q_{\mathcal{I}}$ take the shape x_a with $a \in \Delta^{\mathcal{I}}$, and $q_{\mathcal{I}}$ contains atoms $A(x_a)$ for all $a \in \Delta^{\mathcal{I}}$ and $r(x_a, x_b)$ for all $(a, b) \in r^{\mathcal{I}}$. Answer variables of $q_{\mathcal{I}}$ are all x_a with $a \in \mathbf{a}$, the remaining variables are quantified. Then it can be verified that $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}} \models q_{\mathcal{I}}(\mathbf{a})$ and $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}} \not\models q_{\mathcal{I}}(\mathbf{a})$. \square

We now show that it is possible to refine Lemma 16 so that it makes a much more careful statement in which bounded homomorphisms and unbounded ones are mixed. It is then possible to check the unbounded homomorphism part of the characterization using tree automata as desired, and to deal with bounded homomorphisms using a mosaic technique that “precompiles” relevant information to be used in the automaton construction.

We begin with a useful lemma which shows that, under certain additional conditions, we can extract an unbounded homomorphism from a suitable family of bounded ones. Let \mathcal{I} be an interpretation. For a signature \mathbf{S} , we say that the interpretation \mathcal{I} is **S**-connected if between any two domain elements there is a path using only **S**-role edges, and *finitely branching* if every domain element has only finitely many directly connected elements.

Lemma 17 *Let $\mathcal{I}_1, \mathcal{I}_2$ be finitely branching interpretations and let \mathcal{I}_1 be **S**-connected, for a signature **S**. If there are $d_0 \in \Delta^{\mathcal{I}_1}$ and $e_0 \in \Delta^{\mathcal{I}_2}$ such that for each $i \geq 0$ there is an **S**-homomorphism h_i from $\mathcal{I}_1|_{d_0^i}$ to \mathcal{I}_2 with $h_i(d_0) = e_0$, then $\mathcal{I}_1 \rightarrow_{\mathbf{S}} \mathcal{I}_2$.*

Proof. We are going to construct an \mathbf{S} -homomorphism h from \mathcal{I}_1 to \mathcal{I}_2 step by step, obtaining the desired homomorphism in the limit. We will take care that, at all times, the domain of h is finite and

- (*) there is a sequence h_0, h_1, \dots with h_i an \mathbf{S} -homomorphism from $\mathcal{I}_1|_i^{d_0}$ to \mathcal{I}_2 such that whenever $h(d)$ is already defined, then $h_i(d) = h(d)$ for all $i \geq 0$.

Start with setting $h(d_0) = e_0$. The original sequence h_0, h_1, \dots from the lemma witnesses (*). Now consider the set Λ that consists of all elements $d \in \Delta^{\mathcal{I}_1}$ such that $h(d)$ is undefined and there is an $e \in \Delta^{\mathcal{I}_1}$ with $h(e)$ defined and $(e, d) \in r^{\mathcal{I}_1}$, for some \mathbf{S} -role r . Since the domain of h is finite and \mathcal{I}_1 is finitely branching, Λ is finite. By (*) and since \mathcal{I}_2 is finitely branching, for each $d \in \Lambda$, there are only finitely many e' such that $h_i(d) = e'$ for some i . Thus, there must be a function $\delta : \Lambda \rightarrow \Delta^{\mathcal{I}_2}$ such that, for infinitely many i , we have $h_i(d) = \delta(d)$, for all $d \in \Lambda$. Extend h accordingly, that is, set $h(d) = \delta(d)$ for all $d \in \Lambda$. Clearly, the sequence h_0, h_1, \dots from (*) before the extension is no longer sufficient to witness (*) after the extension. We fix this by skipping homomorphisms that do not respect δ , that is, define a new sequence h'_0, h'_1, \dots by using as h'_i the restriction of h_j to the domain of $\mathcal{I}_1|_i^{d_0}$ where $j \geq i$ is smallest such that $h_j(d) = \delta(d)$ for all $d \in \Lambda$. This finishes the construction. The lemma follows from the fact that, due to \mathbf{S} -connectedness of \mathcal{I}_1 , every element is eventually reached. Note that we automatically have $h(a) = a$ for all individual names a (as required), no matter whether d_0 is an individual name or not. \square

We continue with introducing notation relevant for the refinement of Lemma 16. For a signature \mathbf{S} , we use $\mathcal{I}|_{\mathbf{S}}^{\text{con}}$ to denote the restriction of the interpretation \mathcal{I} to those elements that can be reached from an ABox individual by traveling along \mathbf{S} -roles (forwards or backwards). An \mathbf{S} -subtree in $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ is a maximal weakly tree-shaped, \mathbf{S} -connected sub-interpretation \mathcal{I} of $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$ that does not comprise any ABox individuals. The *root* of \mathcal{I} is the (unique) element of $\Delta^{\mathcal{I}}$ that can be reached from an ABox individual on a shortest path among all elements of $\Delta^{\mathcal{I}}$.

Theorem 18 *Let \mathcal{T}_1 and \mathcal{T}_2 be Horn-ALCHIF TBoxes with $\mathcal{T}_1 \models_{\mathbf{A}, \mathbf{Q}}^{\text{RI}} \mathcal{T}_2$. Then $\mathcal{T}_1 \models_{\mathbf{A}, \mathbf{Q}}^{\text{CQ}} \mathcal{T}_2$ iff for all tree-shaped \mathbf{A} -ABoxes \mathcal{A} consistent with \mathcal{T}_1 and \mathcal{T}_2 , and for all weakly tree-shaped and finitely branching models \mathcal{I}_1 of \mathcal{A} and \mathcal{T}_1 , the following hold:*

- (1) $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}|_{\mathbf{Q}}^{\text{con}} \rightarrow_{\mathbf{Q}} \mathcal{I}_1$;
- (2) for all \mathbf{Q} -subtrees \mathcal{I} in $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$, one of the following holds:
 - (a) $\mathcal{I} \rightarrow_{\mathbf{Q}} \mathcal{I}_1$;
 - (b) $\mathcal{I} \rightarrow_{\mathbf{Q}}^{\text{fin}} \mathcal{I}_{\mathcal{T}_1, \text{tp}_{\mathcal{I}_1}(a)}$ for some $a \in \text{ind}(\mathcal{A})$.

Proof. “ \Leftarrow ”. Suppose that $\mathcal{T}_1 \not\models_{\mathbf{A}, \mathbf{Q}}^{\text{CQ}} \mathcal{T}_2$. By Proposition 13, there is a tree-shaped \mathbf{A} -ABox \mathcal{A} consistent with both \mathcal{T}_i , and a weakly tree-shaped \mathbf{Q} -CQ q such that either

- (1') q has a single answer variable and there is an element $a \in \text{ind}(\mathcal{A})$ such that $\mathcal{T}_2, \mathcal{A} \models q(a)$ but $\mathcal{T}_1, \mathcal{A} \not\models q(a)$, or
- (2') q is Boolean and $\mathcal{T}_2, \mathcal{A} \models q$ but $\mathcal{T}_1, \mathcal{A} \not\models q$.

In case (1'), let π be a match of q in $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$ with $\pi(x) = a$. Since q contains an answer variable, we must have $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}|_{\mathbf{Q}}^{\text{con}} \not\rightarrow_{\mathbf{Q}} \mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$ as otherwise the composition of π and the witnessing homomorphism shows $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}} \models q(a)$, which is not the case. Thus, Condition (1) is violated for $\mathcal{I}_1 = \mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$.

In case (2'), consider a match π of q in $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$, and let \mathcal{I} be the restriction of $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$ to the elements in the range of π . Clearly, we have $\mathcal{I} \not\rightarrow_{\mathbf{Q}} \mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$. Consequently, we also have $\mathcal{I} \not\rightarrow_{\mathbf{Q}}^n \mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$ where n is the number of variables in q , implying that Conditions (2a) and (2b) are both false.

“ \Rightarrow ”. Assume that $\mathcal{T}_1 \models_{\mathbf{A}, \mathbf{Q}}^{\text{CQ}} \mathcal{T}_2$ and let \mathcal{A} be a tree-shaped \mathbf{A} -ABox consistent with both \mathcal{T}_i . Moreover, let \mathcal{I}_1 be a weakly tree-shaped, finitely branching model of \mathcal{A} and \mathcal{T}_1 .

For Condition (1), and let $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}|_{\mathbf{Q}}^{\text{con}}$ be the disjoint union of the connected interpretations $\mathcal{J}_1, \dots, \mathcal{J}_k$. In each \mathcal{J}_i , we find at least one individual a_i from $\text{ind}(\mathcal{A})$. Let $\ell \in \{1, \dots, k\}$. By $\mathcal{T}_1 \models_{\mathbf{A}, \mathbf{Q}}^{\text{CQ}} \mathcal{T}_2$ and Lemma 16, we have $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}} \rightarrow_{\mathbf{Q}}^{\text{fin}} \mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$ and hence, by Lemma 11 (2), $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}} \rightarrow_{\mathbf{Q}}^{\text{fin}} \mathcal{I}_1$. By Lemma 10 (3), we find a sequence h_0, h_1, \dots such that h_i is a \mathbf{Q} -homomorphism from $\mathcal{J}_\ell|_i^{a_\ell}$ to \mathcal{I}_1 . Note that we must have $h_i(a_\ell) = a_\ell$ for all i . Thus, Lemma 17 yields $\mathcal{J}_\ell \rightarrow_{\mathbf{Q}} \mathcal{I}_1$ and, in summary, $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}|_{\mathbf{Q}}^{\text{con}} \rightarrow_{\mathbf{Q}} \mathcal{I}_1$.

For Condition (2), let \mathcal{I} be a \mathbf{Q} -subtree in $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$ with root d_0 . By Lemma 16, we have $\mathcal{I} \rightarrow_{\mathbf{Q}}^{\text{fin}} \mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$. By Lemma 10 (3), there is a sequence h_0, h_1, \dots such that h_i is a \mathbf{Q} -homomorphism from $\mathcal{I}|_i^{d_0}$ to $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$. We distinguish two cases, leading to (2)(a) and (2)(b), respectively.

First assume that there is an $e_0 \in \Delta^{\mathcal{I}_{\mathcal{T}_1, \mathcal{A}}}$ such that $h_i(d_0) = e_0$ for infinitely many i . Construct a new sequence h'_0, h'_1, \dots with h'_i a \mathbf{Q} -homomorphism from $\mathcal{I}|_i^{d_0}$ to $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$ by skipping homomorphisms that do not map d_0 to e_0 , that is, h'_i is the restriction of h_j to the domain of $\mathcal{I}|_i^{d_0}$ where $j \geq i$ is smallest such that $h_j(d_0) = e_0$. Clearly, $h'_i(d_0) = e_0$ for all i . Thus, Lemma 17 yields $\mathcal{I} \rightarrow_{\mathbf{Q}} \mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$ and thus, by Lemma 11 (2) $\mathcal{I} \rightarrow_{\mathbf{Q}} \mathcal{I}_1$.

Otherwise, there is no $e_0 \in \Delta^{\mathcal{I}_{\mathcal{T}_1, \mathcal{A}}}$ such that $h_i(d_0) = e_0$ for infinitely many i . We can assume that there is an $a_0 \in \text{ind}(\mathcal{A})$ such that $h_i(d_0) \in \Delta^{\mathcal{I}_{\mathcal{T}_1, \mathcal{A}}|_{a_0}}$ for all i ; in fact, there must be an a_0 such that $h_i(d_0) \in \Delta^{\mathcal{I}_{\mathcal{T}_1, \mathcal{A}}|_{a_0}}$ for infinitely many i and we can again skip homomorphisms to achieve this for all i . It is important to note that the remaining homomorphisms do not necessarily all map domain elements of \mathcal{I} to elements in $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}}|_{a_0}$ due to the presence of inverse roles. Now, since $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$ is finitely branching, for all $i, n \geq 0$ we must find a $j \geq i$ such that $h_j(d_0)$ is a domain element whose distance from a_0 exceeds n (otherwise the previous case would apply). Based on this observation, we construct a sequence of homomorphisms h'_0, h'_1, \dots as follows. For all $i \geq 0$, let h'_i be the restriction of some h'_j to the domain of $\mathcal{I}|_i^{d_0}$ where $j \geq i$ is smallest such that the distance of $h_j(d_0)$ from a_0 exceeds i . It should be clear that each h'_i is a \mathbf{Q} -homomorphism from $\mathcal{I}|_i^{d_0}$ to $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}}|_{a_0}$, hence $\mathcal{I} \rightarrow_{\mathbf{Q}}^{\text{fin}} \mathcal{I}_{\mathcal{T}_1, \mathcal{A}}|_{a_0}$, by Lemma 10 (3).

Now Lemma 11 (2) and Condition (1) of homomorphisms imply $\text{tp}_{\mathcal{I}_{\mathcal{T}_1, \mathcal{A}}}(a_0) \subseteq \text{tp}_{\mathcal{I}_1}(a_0)$; thus by Lemma 11 (4) we have $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}}|_{a_0} \rightarrow \mathcal{I}_{\mathcal{T}_1, \text{tp}_{\mathcal{I}_1}(a_0)}$. Altogether we get $\mathcal{I} \rightarrow_{\mathbf{Q}}^{\text{fin}} \mathcal{I}_{\mathcal{T}_1, \text{tp}_{\mathcal{I}_1}(a_0)}$ as required. \square

3.3 Characterization for 1tCQ Entailment

We now give a 1tCQ version of Theorem 18, that is, a characterization of 1tCQ-entailment. We use (unbounded) simulations in place of homomorphisms.

Theorem 19 *Let \mathcal{T}_1 and \mathcal{T}_2 be Horn-ALCHIF TBoxes with $\mathcal{T}_1 \models_{\mathbf{A}, \mathbf{Q}}^{\text{RI}} \mathcal{T}_2$. Then the following three conditions are equivalent:*

- (a) $\mathcal{T}_1 \models_{\mathbf{A}, \mathbf{Q}}^{\text{1tCQ}} \mathcal{T}_2$;
- (b) $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}} \preceq_{\mathbf{Q}} \mathcal{I}_1$ for all tree-shaped \mathbf{A} -ABoxes \mathcal{A} that are consistent with \mathcal{T}_1 and \mathcal{T}_2 , and for all weakly tree-shaped, finitely branching models \mathcal{I}_1 of \mathcal{A} and \mathcal{T}_1 .

Proof.

(b) \Rightarrow (a). This implication is analogous to the “ \Rightarrow ” direction in the proof of Theorem 18, but using Proposition 14 in place of Proposition 13. This rules out Case (2') and thus Condition (2) from Theorem 18. In Condition (1), we can replace homomorphisms with simulations since the witnessing CQ is a 1tCQ.

(a) \Rightarrow (b). For the proof of this direction, we need a bounded variant of simulations, analogously to bounded homomorphisms. We write $\mathcal{I}_1 \preceq_{\mathbf{Q}}^n \mathcal{I}_2$ if for any subinterpretation \mathcal{I}'_1 of \mathcal{I}_1 with $|\Delta^{\mathcal{I}'_1}| \leq n$, we have $\mathcal{I}'_1 \preceq_{\mathbf{Q}} \mathcal{I}_2$. Moreover, we write $\mathcal{I}_1 \preceq_{\mathbf{Q}}^{\text{fin}} \mathcal{I}_2$ if $\mathcal{I}_1 \preceq_{\mathbf{Q}}^n \mathcal{I}_2$, for any n .

Assume that $\mathcal{T}_1 \models_{\mathbf{A}, \mathbf{Q}}^{\text{1tCQ}} \mathcal{T}_2$ and let \mathcal{A} be a tree-shaped \mathbf{Q} -ABox that is consistent with both \mathcal{T}_i . Let \mathcal{I}_1 be a weakly tree-shaped, finitely branching model of \mathcal{T}_1 and \mathcal{A} . We first show the following.

Claim. $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}|_a \preceq_{\mathbf{Q}}^{\text{fin}} \mathcal{I}_1$ for all $a \in \text{ind}(\mathcal{A})$.

Assume to the contrary that $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}|_a \not\preceq_{\mathbf{Q}}^{\text{fin}} \mathcal{I}_1$. Then $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}|_a \not\preceq_{\mathbf{Q}}^n \mathcal{I}_1$ for some n , that is, there is a subinterpretation \mathcal{I} of $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}|_a$ with $|\Delta^{\mathcal{I}}| \leq n$ such that $\mathcal{I} \not\preceq_{\mathbf{Q}} \mathcal{I}_1$. We can assume w.l.o.g. that \mathcal{I} is connected and contains a (otherwise we just extend \mathcal{I} and increase n accordingly). Let $q_{\mathcal{I}}$ be \mathcal{I} viewed as a weakly tree-shaped CQ with domain elements viewed as variables and a viewed as the only answer variable. Clearly, $\mathcal{I} \models q(a)$ and thus $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}|_a \models q(a)$. Let π be a witnessing match of q in \mathcal{I} . To transform q into a 1tCQ, perform the following operations:

- remove all atoms of the form $r(a, a)$;
- Split multi-edges in q by duplicating subtree queries: if $r_1(z, z'), \dots, r_n(z, z')$ are in q where r_1, \dots, r_n are potentially inverse roles and z' is a successor of z when the answer variable is viewed as the root of q , then introduce n copies of the subtree rooted at z' and connect z to each copy with one of the roles r_1, \dots, r_n . Remove the subtree rooted at z' .

The result of this transformation is a 1tCQ q' , which still satisfies $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}} \models q'(a)$. On the other hand, $\mathcal{I}_1 \not\models q'(a)$ because, otherwise, a match π' of $q'(x)$ in \mathcal{I}_1 would give rise to a simulation of \mathcal{I} in \mathcal{I}_1 . This finishes the proof of the claim.

The claim implies that for all $a \in \text{ind}(\mathcal{A})$ and all $i \geq 0$, there is a \mathbf{Q} -simulation from $(\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}|_a)|_i^a$ to \mathcal{I}_1 : for every a , the situation parallels Condition (3a) of Lemma 10 with

homomorphisms replaced by simulations. We can use exactly the same arguments as in the proof of Lemma 10 (3) to show that there is a \mathbf{Q} -simulation from $(\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}|_a)|_i^a$ to \mathcal{I}_1 . The union of these simulations, for all $a \in \text{ind}(\mathcal{A})$, is the desired \mathbf{Q} -simulation from $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$ to \mathcal{I}_1 . \square

4. Decidability and Complexity

Our aim is to prove the following result.

Theorem 20 *In Horn-ALCHIF and any of its fragments that contains \mathcal{ELI} or Horn-ALC, the following problems are 2EXPTIME-complete: (\mathbf{A}, \mathbf{Q}) -CQ entailment, (\mathbf{A}, \mathbf{Q}) -CQ inseparability, and (\mathbf{A}, \mathbf{Q}) -CQ conservative extensions. This holds even when $\mathbf{A} = \mathbf{Q}$. Moreover, these problems can be solved in time $2^{2^{p(|\mathcal{T}_2| \log |\mathcal{T}_1|)}}$, where p is a polynomial.*

The lower bounds have been established by Botoeva et al. (2019). To obtain the upper bounds, we use an approach based on the characterization provided by Theorem 18 that combines a mosaic technique with tree automata. More precisely, we first observe that Condition (2b) is the only part of Theorem 18 concerned with bounded homomorphisms, and that addressing this condition requires us to check for a given \mathbf{Q} -subtree \mathcal{I} in $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$ and a \mathcal{T}_1 -type t_1 , whether $\mathcal{I} \rightarrow_{\mathbf{Q}}^{\text{fin}} \mathcal{I}_{\mathcal{T}_1, t_1}$. In fact, it suffices to consider subtrees \mathcal{I} whose root is connected only by non- \mathbf{Q} roles as all \mathcal{I} that do not satisfy this condition are not maximal or treated by Condition (1) of Theorem 18. An analysis of universal models reveals that such \mathcal{I} are completely determined by the \mathcal{T}_2 -type t_2 of their root and are in fact identical to $\mathcal{I}_{\mathcal{T}_2, t_2}|_{\mathbf{Q}}^{\text{con}}$.⁴ This gives rise to the problem of deciding whether $\mathcal{I}_{\mathcal{T}_2, t_2}|_{\mathbf{Q}}^{\text{con}} \rightarrow_{\mathbf{Q}}^{\text{fin}} \mathcal{I}_{\mathcal{T}_1, t_1}$, for given $\mathcal{T}_1, \mathcal{T}_2, t_1, t_2$, and \mathbf{Q} . We show that this problem can be solved in 2EXPTIME using a mosaic technique and then use that result as a black box in an automata based approach to prove the upper bound.

4.1 Mosaics

The aim of this section is to prove the following.

Theorem 21 *Given two Horn-ALCHIF TBoxes \mathcal{T}_1 and \mathcal{T}_2 , a \mathcal{T}_1 -type t_1 , a \mathcal{T}_2 -type t_2 , and a signature \mathbf{S} , it can be decided in time $2^{2^{p(|\mathcal{T}_2| \log |\mathcal{T}_1|)}}$ whether $\mathcal{I}_{\mathcal{T}_2, t_2}|_{\mathbf{S}}^{\text{con}} \rightarrow_{\mathbf{S}}^{\text{fin}} \mathcal{I}_{\mathcal{T}_1, t_1}$, where p is a polynomial.*

We prove Theorem 21 by replacing bounded homomorphisms with unbounded ones and, to compensate for this, further replacing the target interpretation $\mathcal{I}_{\mathcal{T}_1, t_1}$ with a suitably constructed class of interpretations.

To illustrate, consider Example 15 and let $t_1 = t_2 = \{B\}$. The difference between $\mathcal{I}_{\mathcal{T}_2, t_2} \rightarrow_{\mathbf{S}}^{\text{fin}} \mathcal{I}_{\mathcal{T}_1, t_1}$ and $\mathcal{I}_{\mathcal{T}_2, t_2} \rightarrow_{\mathbf{S}} \mathcal{I}_{\mathcal{T}_1, t_1}$ is that unbounded homomorphisms fail once they “reach the root” of $\mathcal{I}_{\mathcal{T}_1, t_1}$ while bounded homomorphisms can, depending on the bound, map the root of $\mathcal{I}_{\mathcal{T}_2, t_2}$ deeper and deeper into $\mathcal{I}_{\mathcal{T}_1, t_1}$, thus never reaching its root. The latter is possible because $\mathcal{I}_{\mathcal{T}_1, t_1}$ is regular in the sense that any two elements which have the same type are the root of isomorphic subtree interpretations. This is of course not only true in

4. This depends on the assumption that the root is connected only by non- \mathbf{Q} roles, c.f. Condition (ii) from the definition of universal models.

this example, but in any universal model. To transition back from bounded to unbounded homomorphisms, we replace $\mathcal{I}_{\mathcal{T}_1, t_1}$ with a class of interpretations that can be seen as a “backwards regularization” of $\mathcal{I}_{\mathcal{T}_1, t_1}$. In our concrete example, $\mathcal{I}_{\mathcal{T}_1, t_1}$ contains an element that has the same type as the root of $\mathcal{I}_{\mathcal{T}_1, t_1}$ and that has a predecessor d . This would lead us to consider (among others) the interpretation obtained from $\mathcal{I}_{\mathcal{T}_1, t_1}$ by adding a predecessor of the same type as d and possibly adding further predecessors, even ad infinitum. We now make this idea precise.

A *rooted weakly tree-shaped interpretation* is a pair (\mathcal{I}, d_0) with \mathcal{I} a weakly tree-shaped interpretation and $d_0 \in \Delta^{\mathcal{I}}$ a *root*. The root imposes a direction on \mathcal{I} and thus allows us to speak about successors and predecessors, that is, when $(d, e) \in r^{\mathcal{I}}$, then e is a *successor* of d if the distance of e to d_0 in the undirected graph $(\Delta^{\mathcal{I}}, \{\{d, e\} \mid (d, e) \in r^{\mathcal{I}} \text{ for some } r \in \mathbf{N}_{\mathcal{C}}\})$ is larger than the distance of d to d_0 , and a *predecessor* otherwise. Note that the direction does not correspond to the distinction between roles and inverse roles, but rather reflects the directedness of universal models $\mathcal{I}_{\mathcal{T}, t}$, which can be viewed as being rooted weakly tree-shaped: $\Delta^{\mathcal{I}_{\mathcal{T}, t}}$ is a set of paths and the root of $\mathcal{I}_{\mathcal{T}, t}$ is the unique path in $\Delta^{\mathcal{I}_{\mathcal{T}, t}}$ of minimum length.

Throughout the section, we work with sets ρ of (possibly inverse) roles; in particular, we say that b is a ρ -*successor* of a if $\rho = \{r \mid (a, b) \in r^{\mathcal{I}}\}$. Let \mathcal{T} be a Horn- \mathcal{ALCHIF} TBox and let $\mathbf{tp}(\mathcal{T})$ be the set of all \mathcal{T} -types consistent with \mathcal{T} . For every $t_0 \in \mathbf{tp}(\mathcal{T})$, we use $\mathbf{tp}(\mathcal{T}, t_0)$ to denote the set of all $t \in \mathbf{tp}(\mathcal{T})$ that occur in the universal model $\mathcal{I}_{\mathcal{T}, t_0}$ of t_0 and \mathcal{T} . Furthermore, given a rooted weakly tree-shaped interpretation (\mathcal{I}, d_0) and an element $d \in \Delta^{\mathcal{I}}$, the *1-neighborhood of d in \mathcal{I}* is a tuple $n_1^{\mathcal{I}}(d) = (t^-, \rho, t, S)$ such that

- (a) $t = \mathbf{tp}_{\mathcal{I}}(d)$;
- (b) if d is the ρ' -successor of some $d^- \in \Delta^{\mathcal{I}}$, then $t^- = \mathbf{tp}_{\mathcal{I}}(d^-)$ and $\rho = \rho'$, otherwise $\rho = t^- = \perp$;
- (c) S is the set of all pairs (ρ', t') such that there is a ρ' -successor d' of d with $t' = \mathbf{tp}_{\mathcal{I}}(d')$.

We further write $(t_1^-, \rho_1, t_1, S_1) \sqsubseteq (t_2^-, \rho_2, t_2, S_2)$ if $t_1 = t_2$, $S_1 \subseteq S_2$ and, if $\rho_1 \neq \perp$, then $\rho_1 = \rho_2$ and $t_1^- = t_2^-$.

In the following, we define the class $\mathbf{mod}_{\omega}(\mathcal{T}, t_0)$ of models of \mathcal{T} that we consider as homomorphism targets. Each model from this class is obtained as the limit of a (finite or infinite) sequence of rooted weakly tree-shaped interpretations. Start to construct the sequence by choosing a type $t \in \mathbf{tp}(\mathcal{T}, t_0)$ and defining $\mathcal{I} = (\{d_0\}, \cdot^{\mathcal{I}})$ such that $\mathbf{tp}_{\mathcal{I}}(d_0) = t$. The remaining sequence is obtained by repeatedly applying the following rule:

- (R) Let $d \in \Delta^{\mathcal{I}}$. Choose some $e \in \Delta^{\mathcal{I}_{\mathcal{T}, t_0}}$ such that $n_1^{\mathcal{I}}(d) \sqsubseteq n_1^{\mathcal{I}_{\mathcal{T}, t_0}}(e)$, and add to d the predecessor and/or successors required to achieve $n_1^{\mathcal{I}}(d) = n_1^{\mathcal{I}_{\mathcal{T}, t_0}}(e)$. If a new predecessor is added, it becomes the new root.

Note that rule application does not need to be fair nor exhaustive. Now, $\mathbf{mod}_{\omega}(\mathcal{T}, t_0)$ is the set of all (weakly tree-shaped but not rooted) interpretations \mathcal{I} that can be obtained as the limit of a sequence constructed in the described way.

Lemma 22 *Let \mathcal{T} be a Horn- \mathcal{ALCHIF} TBox, $t_0 \in \mathbf{tp}(\mathcal{T})$, and \mathcal{I} a finitely branching, weakly tree-shaped interpretation. Then $\mathcal{I} \rightarrow_{\mathbf{S}}^{\text{fin}} \mathcal{I}_{\mathcal{T}, t_0}$ iff there is a $\mathcal{J} \in \mathbf{mod}_{\omega}(\mathcal{T}, t_0)$ with $\mathcal{I} \rightarrow_{\mathbf{S}} \mathcal{J}$.*

Proof. “ \Rightarrow ”. Suppose $\mathcal{I} \rightarrow_{\mathbf{S}}^{\text{fin}} \mathcal{I}_{\mathcal{T},t_0}$ for some finitely branching, weakly tree-shaped interpretation \mathcal{I} , and choose a root d_0 of \mathcal{I} . By Lemma 10 (3), there is a sequence h_0, h_1, \dots such that h_i is an \mathbf{S} -homomorphism from $\mathcal{I}|_i^{d_0}$ to $\mathcal{I}_{\mathcal{T},t_0}$. Note that both \mathcal{I} and $\mathcal{I}_{\mathcal{T},t_0}$ are finitely branching. By skipping homomorphisms (similar to the proof of Lemma 17), we can thus construct a new sequence h'_0, h'_1, \dots such that h'_i is an \mathbf{S} -homomorphism from $\mathcal{I}|_i^{d_0}$ to $\mathcal{I}_{\mathcal{T},t_0}$ and, additionally, for every $i \geq 0$ and every $d \in \Delta^{\mathcal{I}|_i^{d_0}}$ the following properties hold:

- (i) $n_1^{\mathcal{I}_{\mathcal{T},t_0}}(h'_i(d)) = n_1^{\mathcal{I}_{\mathcal{T},t_0}}(h'_j(d))$ for all j with $j \geq i$;
- (ii) if e is a successor of d in \mathcal{I} , then one of the following is the case:
 - $h'_j(e)$ is the predecessor of $h'_j(d)$ in $\mathcal{I}_{\mathcal{T},t_0}$, for all j with $j \geq i$, or
 - $h'_j(e)$ is a successor of $h'_j(d)$ in $\mathcal{I}_{\mathcal{T},t_0}$, and there is some (ρ, t) in component S of $n_1^{\mathcal{I}_{\mathcal{T},t_0}}(h'_i(d))$ such that for all j with $j \geq i$, we have $\rho = \{r \mid (h'_j(d), h'_j(e)) \in r^{\mathcal{I}_{\mathcal{T},t_0}}\}$ and $h'_j(e) = t$.

Guided by h'_i , we construct a sequence of rooted weakly tree-shaped interpretations $(\mathcal{J}_0, e_0), (\mathcal{J}_1, e_1), \dots$ and a sequence g_0, g_1, \dots with g_i an \mathbf{S} -homomorphism from $\mathcal{I}|_i^{d_0}$ to \mathcal{J}_i such that for every i, j with $0 \leq i \leq j$ and every $d \in \Delta^{\mathcal{I}|_i^{d_0}}$, we have $g_i(d) = g_j(d)$. The required interpretation $\mathcal{J} \in \text{mod}_{\omega}(\mathcal{T}, t_0)$ is obtained in the limit. Throughout the construction, we maintain the invariant

$$n_1^{\mathcal{J}_i}(g_i(d)) \sqsubseteq n_1^{\mathcal{I}_{\mathcal{T},t_0}}(h'_i(d)) \tag{*}$$

for all i, d such that $g_i(d)$ is defined.

We start with $\mathcal{J}_0 = (\{e_0\}, \cdot^{\mathcal{J}_0})$ such that $\text{tp}_{\mathcal{J}_0}(e_0) = \text{tp}_{\mathcal{I}_{\mathcal{T},t_0}}(h'_0(d_0))$, choose e_0 as the root and set $g_0(d_0) = e_0$. Clearly (*) is satisfied. Assuming that (\mathcal{J}_i, e_i) and g_i are already defined, we extend them to $(\mathcal{J}_{i+1}, e_{i+1})$ and g_{i+1} as follows. Choose some $d \in \Delta^{\mathcal{I}|_i^{d_0}}$ and $d' \notin \Delta^{\mathcal{I}|_i^{d_0}}$ such that $(d, d') \in r^{\mathcal{I}}$ for some role r . By invariant (*) and Point (i), we have $n_1^{\mathcal{J}_i}(g_i(d)) \sqsubseteq n_1^{\mathcal{I}_{\mathcal{T},t_0}}(h'_j(d))$ for all $j \geq i$. Thus, we can apply **(R)** to $g_i(d)$ in \mathcal{J}_i and $h'_i(d)$ in $\Delta^{\mathcal{I}_{\mathcal{T},t_0}}$. More precisely, we obtain \mathcal{J}_{i+1} from \mathcal{J}_i by adding a predecessor and/or successors to achieve

$$n_1^{\mathcal{J}_{i+1}}(g_i(d)) = n_1^{\mathcal{I}_{\mathcal{T},t_0}}(h'_i(d)). \tag{**}$$

Moreover, e_{i+1} is the root of \mathcal{J}_{i+1} , updated according to **(R)**. To define g_{i+1} , we extend g_i to d' by distinguishing two cases according to Point (ii):

- Suppose $h'_j(d')$ is the predecessor of $h'_j(d)$ for all $j \geq i$. We set $g_{i+1}(d')$ to the predecessor of $g_i(d)$ (exists due to (**)). Clearly, (*) is satisfied also for $g_{i+1}(d')$.
- Suppose $h'_j(d')$ is a successor of $h'_j(d)$ for all $j \geq i$, and let (ρ, t) be as described. Let $e \in \Delta^{\mathcal{J}_{i+1}}$ be an element with $\text{tp}_{\mathcal{J}_{i+1}}(e) = t$ and $\rho = \{r \mid (g_i(d), e) \in r^{\mathcal{J}_{i+1}}\}$. Note that such an element exists due to (**). Set $g_{i+1}(d') = e$. Clearly, (*) is satisfied also for $g_{i+1}(d')$.

The construction of \mathcal{J} and h is finished by setting $h = \bigcup_{i \geq 0} g_i$ and $\mathcal{J}' = \bigcup_{i \geq 0} \mathcal{J}_i$.

“ \Leftarrow ”. Suppose there is a $\mathcal{J} \in \text{mod}_\omega(\mathcal{T}, t_0)$ with $\mathcal{I} \rightarrow_{\mathbf{S}} \mathcal{J}$. It suffices to show $\mathcal{J} \rightarrow^{\text{fin}} \mathcal{I}_{\mathcal{T}, t_0}$. Let $(\mathcal{J}_0, d_0), (\mathcal{J}_1, d_1), \dots$ be the sequence of rooted weakly tree-shaped interpretations whose limit is \mathcal{J} . We verify the following claim, which implies $\mathcal{J} \rightarrow^{\text{fin}} \mathcal{I}_{\mathcal{T}, t_0}$.

Claim. For all $i \geq 0$, we have:

- (i) there is an $e_0 \in \Delta^{\mathcal{I}_{\mathcal{T}, t_0}}$ with $\text{tp}_{\mathcal{J}_i}(d_i) = \text{tp}_{\mathcal{I}_{\mathcal{T}, t_0}}(e_0)$;
- (ii) for all $e_0 \in \Delta^{\mathcal{I}_{\mathcal{T}, t_0}}$ with $\text{tp}_{\mathcal{J}_i}(d_i) \subseteq \text{tp}_{\mathcal{I}_{\mathcal{T}, t_0}}(e_0)$, we have $(\mathcal{J}_i, d_i) \rightarrow (\mathcal{I}_{\mathcal{T}, t_0}, e_0)$.

We prove the claim by induction on i . For $i = 0$, Points (i) and (ii) follow from the definition of \mathcal{J}_0 . For the inductive step, consider \mathcal{J}_{i+1} and suppose **(R)** has been applied to some $d \in \Delta^{\mathcal{J}_i}$ and $e \in \Delta^{\mathcal{I}_{\mathcal{T}, t_0}}$.

For $i > 0$, observe first that Point (i) is trivially preserved when $d_{i+1} = d_i$. In case d_{i+1} is the predecessor of d_i , it is preserved by the condition on the choice of e in **(R)**: e has the same type as d_i and the predecessor e' of e has the same type as d_{i+1} .

For Point (ii), we distinguish two cases.

Case 1. Suppose the application of **(R)** has not added any predecessors to d . In particular, we then have $d_{i+1} = d_i$. For Point (ii), take any e_0 with $\text{tp}_{\mathcal{J}_i}(d_{i+1}) \subseteq \text{tp}_{\mathcal{I}_{\mathcal{T}, t_0}}(e_0)$. As $d_i = d_{i+1}$, induction hypothesis implies that there is a homomorphism $h : (\mathcal{J}_i, d_{i+1}) \rightarrow (\mathcal{I}_{\mathcal{T}, t_0}, e_0)$. We extend h to the domain of \mathcal{J}_{i+1} by doing the following for each newly added successor d' of d .

Let $\text{tp}_{\mathcal{J}_{i+1}}(d) = t$ and $\text{tp}_{\mathcal{J}_{i+1}}(d') = t'$ and $\rho = \{r \mid (d, d') \in r^{\mathcal{J}_{i+1}}\}$. By the choice of e in **(R)**, e is of type t and has a ρ -successor of type t' . By definition of the universal model, there is some $r \in \rho$ with $t \rightsquigarrow_r^{\mathcal{T}} t'$ and $\rho = \{s \mid \mathcal{T} \models r \sqsubseteq s\}$. Denote with $\hat{t} = \text{tp}_{\mathcal{I}_{\mathcal{T}, t_0}}(h(d))$. The definition of a homomorphism yields $t \subseteq \hat{t}$. Thus, there is $\hat{t}' \supseteq t'$ such that $\hat{t} \rightsquigarrow_r^{\mathcal{T}} \hat{t}'$. By definition of the universal model, $h(d)$ has a ρ -successor of type \hat{t}' or a ρ -predecessor of type \hat{t}'' , for $\hat{t}'' \supseteq \hat{t}'$. We extend h by setting $h(d')$ to that predecessor or successor, respectively.

After the extension, h witnesses $(\mathcal{J}_{i+1}, d_{i+1}) \rightarrow (\mathcal{I}_{\mathcal{T}, t_0}, e_0)$.

Case 2. Application of **(R)** has added a predecessor d' to d . Then $d_i = d$ and $d_{i+1} = d'$. Let $t = \text{tp}_{\mathcal{J}_{i+1}}(d)$, $t' = \text{tp}_{\mathcal{J}_{i+1}}(d')$ and $\rho = \{r \mid (d', d) \in r^{\mathcal{J}_{i+1}}\}$. By construction of the universal model, there is $r \in \rho$ with $t' \rightsquigarrow_r^{\mathcal{T}} t$ and $\rho = \{s \mid \mathcal{T} \models r \sqsubseteq s\}$. Let e_0 be as in (ii), that is, $\hat{t}' := \text{tp}_{\mathcal{I}_{\mathcal{T}, t_0}}(e_0) \supseteq t'$. We then have that $\hat{t}' \rightsquigarrow_r \hat{t}$ for some $\hat{t} \supseteq t$. By definition of the universal model, e_0 has a ρ -successor of type \hat{t} or a ρ -predecessor of type $\hat{t}'' \supseteq t$. Let this element be \bar{e}_0 . By induction hypothesis, there is a homomorphism $h : (\mathcal{J}_i, d) \rightarrow (\mathcal{I}_{\mathcal{T}, t_0}, \bar{e}_0)$. We extend h by first setting $h(d') = e_0$ and then extending h to all successors of d as in Case 1.

It is not difficult to verify that h witnesses $(\mathcal{J}_{i+1}, d_{i+1}) \rightarrow (\mathcal{I}_{\mathcal{T}, t_0}, e_0)$. \square

We now use Lemma 22 to develop the mosaic based decision procedure that underlies the proof of Theorem 21. The main idea is that since we cannot compute $\text{mod}_\omega(\mathcal{T}, t_1)$ explicitly, we decompose the interpretations in this class into 1-neighborhoods, of which there are only finitely many. Such a 1-neighborhood is then represented by a mosaic, together with a decoration with sets of \mathcal{T}_2 -types that can be homomorphically embedded into that neighborhood.

Let $\mathcal{T}_1, \mathcal{T}_2$ be Horn-*ALCHIF* TBoxes, t_1 a \mathcal{T}_1 -type, and \mathbf{S} a signature. We denote with $\text{rol}(\mathcal{T}_i)$ the set of all roles r, r^- such that the role name r occurs in \mathcal{T}_i (possibly as an inverse

role). Moreover, for a set of roles ρ and a signature \mathbf{S} , denote with $\rho|_{\mathbf{S}}$ the restriction of ρ to \mathbf{S} -roles, and with ρ^- the set $\{r^- \mid r \in \rho\}$.

Formally, a *mosaic* is a tuple $M = (t^-, \rho, t, S, \ell)$ such that $(t^-, \rho, t, S) = n_1^{\mathcal{I}_{\mathcal{T}_1, t_1}}(d)$ for some $d \in \Delta^{\mathcal{I}_{\mathcal{T}_1, t_1}}$ and $\ell : \{t^-, t\} \cup S \rightarrow 2^{\text{tp}(\mathcal{T}_2)}$ satisfies the following condition:

- (M) For all $\hat{t} \in \ell(t)$ we have $\hat{t} \cap \mathbf{S} \subseteq t$ and, for all $\hat{t}' \in \text{tp}(\mathcal{T}_2)$ and $r \in \text{rol}(\mathcal{T}_2)$ with $\hat{t} \rightsquigarrow_r^{\mathcal{T}_2} \hat{t}'$, one of the following holds for $\sigma = \{s \in \text{rol}(\mathcal{T}_2) \mid \mathcal{T}_2 \models r \sqsubseteq s\}$:
- (a) $\sigma|_{\mathbf{S}} = \emptyset$;
 - (b) $t^- \neq \perp$, for every $s \in \sigma|_{\mathbf{S}}$ we have $s^- \in \rho$, and $\hat{t}' \in \ell(t^-)$;
 - (c) there is $(\rho', t') \in S$ with $\hat{t}' \in \ell(\rho', t')$ and $\sigma|_{\mathbf{S}} \subseteq \rho'$.

To ease notation, we use t_M^- to denote t^- , ρ_M to denote ρ , and likewise for the other components of a mosaic M . We further use \mathcal{M} to denote the set of all mosaics.

We now describe an elimination algorithm in the style of Pratt's type elimination for PDL (Pratt, 1979), but working on mosaics rather than on types; the existence of this algorithm will establish Theorem 21. Let $\mathcal{M}' \subseteq \mathcal{M}$ be a set of mosaics. An $M \in \mathcal{M}'$ is *good in \mathcal{M}'* if the following conditions are satisfied:

1. for each $(\rho, t) \in S_M$, there is an $N \in \mathcal{M}'$ such that $(t_M, \rho, t) = (t_N^-, \rho_N, t_N)$, $\ell_M(\rho, t) = \ell_N(t_N)$, and $\ell_M(t_M) = \ell_N(t_N^-)$;
2. if $t_M^- \neq \perp$, there is $N \in \mathcal{M}'$ with $(\rho_M, t_M) \in S_N$, $t_M^- = t_N$, $\ell_M(t_M^-) = \ell_N(t_N)$, and $\ell_M(t_M) = \ell_N(\rho_M, t_M)$.

Our algorithm computes the sequence $\mathcal{M}_0, \mathcal{M}_1, \dots$ that starts with $\mathcal{M}_0 = \mathcal{M}$ and where \mathcal{M}_{i+1} is obtained from \mathcal{M}_i by removing all mosaics that are not good in \mathcal{M}_i . This sequence eventually stabilizes, say at \mathcal{M}_p . The following lemma establishes the central property of the elimination algorithm.

Lemma 23 *Let $t_2 \in \text{tp}(\mathcal{T}_2)$. Then the following two statements are equivalent:*

1. *there is a $\mathcal{J} \in \text{mod}_\omega(\mathcal{T}_1, t_1)$ such that $\mathcal{I}_{\mathcal{T}_2, t_2}|_{\mathbf{S}^{\text{con}}} \rightarrow_{\mathbf{S}} \mathcal{J}$;*
2. *\mathcal{M}_p contains a mosaic M with $t_2 \in \ell_M(t_M)$.*

Proof. **1 \Rightarrow 2.** Let h be an \mathbf{S} -homomorphism from $\mathcal{I}_{\mathcal{T}_2, t_2}|_{\mathbf{S}^{\text{con}}}$ to some $\mathcal{J} \in \text{mod}_\omega(\mathcal{T}_1, t_1)$. For every $d \in \Delta^{\mathcal{J}}$, denote with $T_h(d)$ the set of all types mapped to d by h , that is,

$$T_h(d) = \{\text{tp}_{\mathcal{I}_{\mathcal{T}_2, t_2}}(e) \mid h(e) = d, e \in \Delta^{\mathcal{I}_{\mathcal{T}_2, t_2}|_{\mathbf{S}^{\text{con}}}}\}.$$

For every element $d \in \Delta^{\mathcal{J}}$, we define a tuple $M(d) = (t^-, \rho, t, S, \ell)$ as follows:

- $(t^-, \rho, t, S) = n_1^{\mathcal{J}}(d)$;
- $\ell(t) = T_h(d)$;
- if there is a predecessor d' of d , then $\ell(t^-) = T_h(d')$; otherwise, set $\ell(t^-) = \emptyset$ (or any other value);

- for every $(\rho, t') \in S$, let d' be the ρ -successor of d with $\text{tp}_{\mathcal{J}}(d') = t'$, and set $\ell(\rho', t') = T_h(d')$.

It is easy to verify that every $M(d)$ obtained in this way is indeed a mosaic. In particular, it follows from the definition of \mathcal{J} that $(t^-, \rho, t, S) = n_1^{\mathcal{I}_{\mathcal{T}_1, t_1}}(d')$ for some $d' \in \Delta^{\mathcal{I}_{\mathcal{T}_1, t_1}}$. Moreover, by the fact that h is a homomorphism, Condition **(M)** is satisfied.

Let $\mathcal{M}(\mathcal{J}) = \{M(d) \mid d \in \Delta^{\mathcal{J}}\}$. It follows from the construction that all mosaics in $\mathcal{M}(\mathcal{J})$ are good in $\mathcal{M}(\mathcal{J})$; hence $\mathcal{M}(\mathcal{J}) \subseteq \mathcal{M}_p$. Finally, let d_0 be the root of $\mathcal{I}_{\mathcal{T}_2, t_2}$ and let $M := M(h(d_0))$. Then $t_2 \in \ell_M(t_M)$ and thus Point 2 of Lemma 23 is satisfied.

2 \Leftarrow **1**. Assume that \mathcal{M}_p contains a mosaic M with $t_2 \in \ell_M(t_M)$. We construct an interpretation \mathcal{J} by stitching together mosaics. Throughout the construction, we maintain a partial function $q : \Delta^{\mathcal{J}} \rightarrow \mathcal{M}_p$ that records for each domain element of \mathcal{J} the mosaic that gave rise to it. We also maintain the following invariant:

$$\text{If } q(d) = (t^-, \rho, t, S, \ell), \text{ then } n_1^{\mathcal{J}}(d) = (t^-, \rho, t, S). \quad (*)$$

We start with defining \mathcal{J} as the interpretation corresponding to the 1-neighborhood represented by M (in the obvious way), and define $q(e_0) = M$, where e_0 is the domain element that corresponds to the center of that 1-neighborhood. By definition, the invariant $(*)$ is satisfied. Then extend \mathcal{J} by doing the following exhaustively in a fair way: Choose some $d \in \mathcal{J}$ such that $q(d)$ is undefined, and

- (b) If d has a predecessor d' such that $q(d') = M'$ then, due to $(*)$, there is $(\rho, t) \in S_{M'}$ such that d is the ρ -successor of d' in \mathcal{J} and $\text{tp}_{\mathcal{J}}(d) = t$. Let $N \in \mathcal{M}_p$ be the mosaic that exists according to Condition 1 of being good for $(\rho, t) \in S_{M'}$. Then extend \mathcal{J} such that $n_1^{\mathcal{J}}(d) = (t_N^-, \rho_N, t_N, S_N)$ and set $q(d) = N$.
- (c) If d has a successor d' such that $q(d') = M'$ then, due to $(*)$, we know that $t_{M'}^- = \text{tp}_{\mathcal{J}}(d) \neq \perp$. Let $N \in \mathcal{M}_p$ be the mosaic that exists according to Condition 2 of being good. Then extend \mathcal{J} such that $n_1^{\mathcal{J}}(d) = (t_N^-, \rho_N, t_N, S_N)$ and set $q(d) = N$.

It is immediate from the construction that this preserves $(*)$, and that one of (a) and (b) always applies. Moreover, by construction, any interpretation \mathcal{J} obtained in the limit of this process is an element of $\text{mod}_{\omega}(\mathcal{T}_1, t_1)$.

It thus remains to construct an **S**-homomorphism h that witnesses $\mathcal{I}_{\mathcal{T}_2, t_2} |_{\mathbf{S}}^{\text{con}} \rightarrow_{\mathbf{S}} \mathcal{J}$. We construct h step by step, essentially following the construction of $\mathcal{I}_{\mathcal{T}_2, t_2}$, maintaining the invariant:

$$\text{If } h(d) \text{ is defined, then } \text{tp}_{\mathcal{I}_{\mathcal{T}_2, t_2}}(d) \in \ell_{q(h(d))}(t_{q(h(d))}). \quad (\dagger)$$

Let d_0 be the root of $\mathcal{I}_{\mathcal{T}_2, t_2}$. We start with setting $h(d_0) = e_0$, where e_0 is as above. By the assumption that $t_2 \in \ell_M(t_M)$, invariant (\dagger) is satisfied. Now, exhaustively apply the following step:

Choose $d \in \Delta^{\mathcal{I}_{\mathcal{T}_2, t_2} |_{\mathbf{S}}^{\text{con}}}$ such that $h(d)$ is not defined but $h(d') = e$ is defined for the predecessor d' of d . Let $t = \text{tp}_{\mathcal{I}_{\mathcal{T}_2, t_2}}(d)$, $t' = \text{tp}_{\mathcal{I}_{\mathcal{T}_2, t_2}}(d')$, and $M' = q(d')$. By definition of $\mathcal{I}_{\mathcal{T}_2, t_2}$, $t' \rightsquigarrow_r^{\mathcal{T}_2} t$ for some $r \in \text{rol}(\mathcal{T}_2)$. Let $\sigma = \{s \mid \mathcal{T} \models r \sqsubseteq s\}$. By Invariant (\dagger) , $t' \in \ell_{M'}(t_{M'})$. Thus, one of (a) to (c) in Condition **(M)** applies. Condition (a) is impossible since $\sigma |_{\mathbf{S}} \neq \emptyset$

as $d, d' \in \Delta^{\mathcal{I}_{\mathcal{T}_2, t_2} | \mathbb{S}^{\text{con}}}$. In case of (b), we extend h by setting $h(d)$ to the predecessor of $h(d')$. In case of (c), we extend h by setting $h(d)$ to the according successor of $h(d')$.

Note that h extended like this satisfies the homomorphism conditions and preserves (\dagger) due to the conditions in (b) and (c). \square

We now finish the proof of Theorem 21. Due to Lemma 22, we can decide $\mathcal{I}_{\mathcal{T}_2, t_2} | \mathbb{S}^{\text{con}} \rightarrow_{\mathbb{S}}^{\text{fin}} \mathcal{I}_{\mathcal{T}_1, t_1}$ by checking whether there is a $\mathcal{J} \in \text{mod}_{\omega}(\mathcal{T}_1, t_1)$ with $\mathcal{I}_{\mathcal{T}_2, t_2} | \mathbb{S}^{\text{con}} \rightarrow_{\mathbb{S}} \mathcal{J}$. By Lemma 23, this can be done by constructing the described sequence of mosaics $\mathcal{M}_0, \mathcal{M}_1, \dots, \mathcal{M}_p$ and checking whether \mathcal{M}_p contains a mosaic M with $t_2 \in T_M$.

The time bound stated in Theorem 21 is a consequence of the following observations. We can compute \mathcal{M}_0 as follows. First, enumerate all possible tuples (t^-, ρ, t, S) with the number of elements in S bounded by $|\mathcal{T}_1|$ (since the outdegree of $\mathcal{I}_{\mathcal{T}_1, t_1}$ is bounded by $|\mathcal{T}_1|$). Then remove those which do not correspond to 1-neighborhoods in $\mathcal{I}_{\mathcal{T}_1, t_1}$. Note that this requires computing \rightsquigarrow_r for every r , which can be done in time $2^{p(|\mathcal{T}_1|)}$, p a polynomial. Moreover, we have to perform reachability checks to verify that the tuple (t^-, ρ, t, S) indeed occurs as the 1-neighborhood in $\mathcal{I}_{\mathcal{T}_1, t_1}$. Based on this set of all tuples, we consider all possible combinations with labelings ℓ and remove those tuples (t^-, ρ, t, S, ℓ) which do not satisfy **(M)**. It is routine to verify that there are at most $2^{2q(|\mathcal{T}_2| \log |\mathcal{T}_1|)}$ many mosaics, for some polynomial q . It remains to note that Conditions 1 and 2 of a mosaic being good can be checked in time polynomial in the size of the current set of mosaics \mathcal{M}_i and that the maximal number of iterations is $|\mathcal{M}_0|$.

4.2 Automata Construction

We now prove the upper bound on CQ entailment stated in Theorem 20 based on the characterization provided by Theorem 18, using the decision procedure asserted by Theorem 21 as a black box. Our main tool are alternating two-way tree automata with counting (2ATA_c) which extend alternating automata on unranked infinite trees as used for example by Grädel and Walukiewicz (1999) with the ability to count.

A *tree* is a non-empty (potentially infinite) set of words $T \subseteq (\mathbb{N} \setminus 0)^*$ closed under prefixes. We assume that trees are finitely branching, that is, for every $w \in T$, the set $\{i > 0 \mid w \cdot i \in T\}$ is finite. For $w \in (\mathbb{N} \setminus 0)^*$, set $w \cdot 0 := w$. For $w = n_0 n_1 \dots n_k$, we set $w \cdot -1 := n_0 \dots n_{k-1}$. For an alphabet Θ , a Θ -labeled tree is a pair (T, L) with T a tree and $L : T \rightarrow \Theta$ a node labeling function.

A 2ATA_c is a tuple $\mathfrak{A} = (Q, \Theta, q_0, \delta, \Omega)$ where Q is a finite set of *states*, Θ is the *input alphabet*, $q_0 \in Q$ is the *initial state*, δ is a *transition function*, and $\Omega : Q \rightarrow \mathbb{N}$ is a *priority function*. The transition function δ maps every state q and input letter $a \in \Theta$ to a positive Boolean formula $\delta(q, a)$ over the truth constants **true** and **false** and *transition atoms* of the form q , \diamond^-q , \square^-q , \diamond_nq and \square_nq . A transition q expresses that a copy of \mathfrak{A} is sent to the current node in state q ; \diamond^-q means that a copy is sent in state q to the predecessor node, which is required to exist; \square^-q means the same except that the predecessor node is not required to exist; \diamond_nq means that a copy of q is sent to n successors and \square_nq means that a copy of q is sent to all but n successors. We use $\diamond q$ and $\square q$ to abbreviate \diamond_1q and \square_0q , respectively. The semantics of 2ATA_c is given in terms of runs as usual, details are in Appendix B. We use $L(\mathfrak{A})$ to denote the set of Θ -labeled trees accepted by \mathfrak{A} . It is not hard to show that 2ATA_c are closed under intersection and that the intersection automaton can be constructed in polynomial time (using techniques from, e.g., Comon, Dauchet, Gilleron,

Löding, Jacquemard, Lugiez, Tison, & Tommasi, 2007). The *emptiness problem* for $2ATA_c$ means to decide, given a $2ATA_c$ \mathfrak{A} , whether $L(\mathfrak{A}) = \emptyset$. We assume here that the numbers in transitions of the form $\diamond_n q$ and $\square_n q$ are encoded in unary.⁵ The following result is obtained via reduction to the emptiness problem of the more standard alternating parity tree automata on ranked trees (Vardi, 1998), see Appendix B.1 for details.

Theorem 24 *The emptiness problem for $2ATA_c$ is in EXPTIME. More precisely, it can be solved in time single exponential in the number of states and the maximal priority, and polynomial in all other components.*

To prove the 2EXPTIME upper bound for CQ entailment stated in Theorem 20, it thus suffices to show that given Horn-*ALCHIF* TBoxes \mathcal{T}_1 and \mathcal{T}_2 and signatures \mathbf{A} and \mathbf{Q} , one can construct in time $2^{2^{p(\mathcal{T}_2|\log|\mathcal{T}_1|)}}$ a $2ATA_c$ \mathfrak{A} with single exponentially many states and both maximal priority and maximal occurring number one such that $L(\mathfrak{A}) \neq \emptyset$ iff $\mathcal{T}_1 \not\models_{\mathbf{A}, \mathbf{Q}}^{\text{CQ}} \mathcal{T}_2$. This is what we do in the following.

The desired $2ATA_c$ runs over Θ -labeled trees with $\Theta = 2^{\Theta_0} \times 2^{\Theta_1} \times 2^{\Theta_2}$ where

- $\Theta_0 = \mathbf{A} \cup \{r^- \mid r \in \mathbf{A}\}$;
- $\Theta_i = \text{sig}(\mathcal{T}_i) \cup \{r^- \mid r \in \text{sig}(\mathcal{T}_i)\}$ for $i = 1, 2$.

For a Θ -labeled tree (T, L) , we use L_i , $i \in \{0, 1, 2\}$ to refer to the i -th component of L , that is, $L(n) = (L_0(n), L_1(n), L_2(n))$, for all $n \in T$. The component L_0 represents a (possibly infinite) ABox

$$\mathcal{A}_L = \{A(n) \mid A \in L_0(n)\} \cup \{r(n \cdot -1, n) \mid n \neq \varepsilon, r \in L_0(n)\},$$

where $r^-(a, b)$ is identified with $r(b, a)$. Note that even when \mathcal{A}_L is finite it is not necessarily a tree-shaped ABox as it might have multi-edges and be disconnected. Components L_1 and L_2 represent interpretations $\mathcal{I}_{L,1} = (T, \cdot^{\mathcal{I}_{L,1}})$ and $\mathcal{I}_{L,2} = (\text{ind}(\mathcal{A}_L), \cdot^{\mathcal{I}_{L,2}})$, where for $i \in \{1, 2\}$:

$$\begin{aligned} A^{\mathcal{I}_{L,i}} &= \{n \mid A \in L_i(n)\} \\ r^{\mathcal{I}_{L,i}} &= \{(n, n \cdot -1) \mid r^- \in L_i(n)\} \cup \{(n \cdot -1, n) \mid r \in L_i(n)\} \end{aligned}$$

Now, let \mathcal{T}_1 and \mathcal{T}_2 be Horn-*ALCHIF* TBoxes and let \mathbf{A} , \mathbf{Q} , and \mathbf{S} be signatures. The claimed $2ATA_c$ \mathfrak{A} is constructed as the intersection of the four $2ATA_c$ s $\mathfrak{A}_1, \mathfrak{A}_2, \mathfrak{A}_3, \mathfrak{A}_4$ provided by the following lemma.

Lemma 25 *There are $2ATA_c$ s $\mathfrak{A}_1, \mathfrak{A}_2, \mathfrak{A}_3, \mathfrak{A}_4$ such that:*

- \mathfrak{A}_1 accepts (T, L) iff \mathcal{A}_L is finite, tree-shaped, and $\varepsilon \in \text{ind}(\mathcal{A}_L)$;
- \mathfrak{A}_2 accepts (T, L) iff $\mathcal{I}_{L,1}$ is a model of \mathcal{A}_L and \mathcal{T}_1 ;
- \mathfrak{A}_3 accepts (T, L) iff \mathcal{A}_L is consistent with \mathcal{T}_2 , and $\mathcal{I}_{L,2}$ is $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}_L}$ restricted to $\text{ind}(\mathcal{A}_L)$;
- \mathfrak{A}_4 accepts (T, L) iff Conditions (1) and (2) from Theorem 18 are not both satisfied, when $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}_L}$ is replaced with $\mathcal{I}_{L,2}$.

5. For binary encoding, Theorem 24 below includes additionally an exponential dependence in the *size* of the largest number appearing in δ .

The number of states of \mathfrak{A}_1 and \mathfrak{A}_2 is polynomial in $|\mathcal{T}_1|$ (and independent of \mathcal{T}_2); the number of states of \mathfrak{A}_3 is polynomial in $|\mathcal{T}_2|$ (and independent of \mathcal{T}_1), and the number of states of \mathfrak{A}_4 is exponential in $|\mathcal{T}_2|$ (and independent of \mathcal{T}_1). All automata can be constructed in time $2^{2^p(|\mathcal{T}_2|^{\log|\mathcal{T}_1|})}$, p a polynomial and have maximal priority of one.

It can be verified that indeed $L(\mathfrak{A}) \neq \emptyset$ iff $\mathcal{T}_1 \not\equiv_{\mathbf{A}, \mathbf{Q}}^{\text{CQ}} \mathcal{T}_2$. The rest of this section is devoted to proving Lemma 25.

AUTOMATA \mathfrak{A}_1 AND \mathfrak{A}_2 .

The construction of the automaton \mathfrak{A}_1 is straightforward and left to the reader. Also \mathfrak{A}_2 is easy to construct. It checks that $\mathcal{I}_{L,1}$ is a model of \mathcal{A}_L by synchronizing the L_0 and L_1 components of the input tree and that all statements in \mathcal{T}_1 are satisfied by $\mathcal{I}_{L,1}$ by imposing constraints on the L_1 component. The latter is particularly simple since \mathcal{T}_1 is in normal form. For the sake of completeness and as a warm up, we present the details. Define $\mathfrak{A}_2 = (Q_2, \Theta, q_0, \delta_2, \Omega_2)$ where

$$Q_2 = \{q_0, q_A\} \cup \{q_\alpha \mid \alpha \in \mathcal{T}_1\} \cup \{q_\rho, \bar{q}_\rho \mid \rho \in \Theta_1\} \cup \{q_{r,B}, q_{r,B}^\downarrow, \bar{q}_{r,B}, \bar{q}_{r,B}^\downarrow \mid \exists r.B \text{ occurs in } \mathcal{T}_1\},$$

and Ω_2 assigns 0 to all states. Here, the transition function δ_2 is given as follows. For $\sigma = (L_0, L_1, L_2)$, set

$$\begin{aligned} \delta_2(q_0, \sigma) &= \Box q_0 \wedge q_A \wedge \bigwedge_{\alpha \in \mathcal{T}_1} q_\alpha \\ \delta_2(q_A, \sigma) &= \bigwedge_{\rho \in L_0} q_\rho \\ \delta_2(q_{\text{func}(r)}, \sigma) &= (q_{r-} \wedge \Box \bar{q}_r) \vee (\bar{q}_{r-} \wedge \Box_1 \bar{q}_r) \\ \delta_2(q_{r \sqsubseteq s}, \sigma) &= \bar{q}_r \vee q_s \\ \delta_2(q_{A_1 \sqcap A_2 \sqsubseteq B}, \sigma) &= \bar{q}_{A_1} \vee \bar{q}_{A_2} \vee q_B \\ \delta_2(q_{A \sqsubseteq \perp}, \sigma) &= \bar{q}_A \\ \delta_2(q_{\top \sqsubseteq A}, \sigma) &= q_A \\ \delta_2(q_{A \sqsubseteq \exists r.B}, \sigma) &= \bar{q}_A \vee q_{r,B} \\ \delta_2(q_{A \sqsubseteq \forall r.B}, \sigma) &= q_B \vee \bar{q}_{r-,A} \\ \delta_2(q_{r,B}, \sigma) &= \Diamond q_{r,B}^\downarrow \vee (q_{r-} \wedge \Diamond^- q_B) \\ \delta_2(\bar{q}_{r,B}, \sigma) &= \Box \bar{q}_{r,B}^\downarrow \wedge (\bar{q}_{r-} \vee \Box^- \bar{q}_B) \\ \delta_2(q_{r,B}^\downarrow, \sigma) &= q_r \wedge q_B \\ \delta_2(\bar{q}_{r,B}^\downarrow, \sigma) &= \bar{q}_r \vee \bar{q}_B. \end{aligned}$$

We further set for all $\rho \in \Theta_1$:

$$\begin{aligned} \delta_2(q_\rho, \sigma) &= \begin{cases} \text{true} & \text{if } \rho \in L_1 \\ \text{false} & \text{if } \rho \notin L_1 \end{cases} \\ \delta_2(\bar{q}_\rho, \sigma) &= \begin{cases} \text{true} & \text{if } \rho \notin L_1 \\ \text{false} & \text{if } \rho \in L_1. \end{cases} \end{aligned}$$

AUTOMATON \mathfrak{A}_4 .

We next consider the automaton \mathfrak{A}_4 as it is the most interesting and crucial ingredient to the construction of \mathfrak{A} . Recall that \mathfrak{A}_4 has to make sure that Conditions (1) and (2) from Theorem 18 are not both satisfied. It achieves this by verifying that its input (T, L) is such that there is an $n \in \text{ind}(\mathcal{A}_L)$ for which one of the following conditions holds, where $\mathcal{I}_{\mathcal{T}_2, L_2(n) \cap \mathbf{N}_C}$ is the universal model of the \mathcal{T}_2 -type $L_2(n) \cap \mathbf{N}_C$ and \mathcal{T}_2 :

1. there is a \mathbf{Q} -role r and an $n' \in \text{ind}(\mathcal{A}_L)$ such that $(n, n') \in r^{\mathcal{I}_{L,2}}$, but $(n, n') \notin r^{\mathcal{I}_{L,1}}$;
2. there is no \mathbf{Q} -homomorphism from $\mathcal{I}_{\mathcal{T}_2, L_2(n) \cap \mathbf{N}_C} |_{\mathbf{Q}}^{\text{con}}$ to $\mathcal{I}_{L,1}$ that maps the root of $\mathcal{I}_{\mathcal{T}_2, L_2(n) \cap \mathbf{N}_C}$ to n ;
3. there is a \mathbf{Q} -subtree \mathcal{I} of $\mathcal{I}_{\mathcal{T}_2, L_2(n) \cap \mathbf{N}_C}$ satisfying the following two conditions.
 - (a) $\mathcal{I} \not\rightarrow_{\mathbf{Q}} \mathcal{I}_{L,1}$;
 - (b) $\mathcal{I} \not\rightarrow_{\mathbf{Q}}^{\text{fin}} \mathcal{I}_{\mathcal{T}_1, \text{tp}_{\mathcal{I}_{L,1}}(m)}$, for all m with $L_0(m) \neq \emptyset$.

For Condition 3, note that every \mathbf{Q} -subtree \mathcal{I} of $\mathcal{I}_{\mathcal{T}_2, L_2(n) \cap \mathbf{N}_C}$ is of the form $\mathcal{I}_{\mathcal{T}_2, t'} |_{\mathbf{Q}}^{\text{con}}$ with $\mathcal{I}_{\mathcal{T}_2, t'}$ the universal model of the \mathcal{T}_2 -type t' satisfied at the root of \mathcal{I} and \mathcal{T}_2 . For a \mathcal{T}_2 -type t , we use $R_{\mathbf{Q}}(t)$ to denote the set of all types that are realized at the root of a \mathbf{Q} -subtree in the universal model $\mathcal{I}_{\mathcal{T}_2, t}$ of t and \mathcal{T}_2 . Then, the \mathbf{Q} -subtrees \mathcal{I} in Condition 3b are exactly the interpretations $\mathcal{I}_{\mathcal{T}_2, t} |_{\mathbf{Q}}^{\text{con}}$, $t \in R_{\mathbf{Q}}(L_2(n) \cap \mathbf{N}_C)$. Also note that Condition 3b is exactly the question addressed by Theorem 21.

The essence of Conditions 2 and 3a is to ensure that there is no \mathbf{Q} -homomorphism from $\mathcal{I}_{\mathcal{T}_2, t} |_{\mathbf{Q}}^{\text{con}}$ to $\mathcal{I}_{L,1}$ that maps the root of $\mathcal{I}_{\mathcal{T}_2, t}$ to n , for some \mathcal{T}_2 -type t and $n \in \text{ind}(\mathcal{A}_L)$ (for Condition 3, one has to consider all $n \in \text{ind}(\mathcal{A}_L)$). The automaton verifies this by making sure that one of the following is true:

- (i) there is a concept name from \mathbf{Q} that is true at the root of $\mathcal{I}_{\mathcal{T}_2, t}$ but not at n in $\mathcal{I}_{L,1}$;
- (ii) there is a \mathcal{T}_2 -type t' and a set of $\text{sig}(\mathcal{T}_2)$ -roles ρ with $t \rightsquigarrow_{\rho}^{\mathcal{T}_2} t'$ such that there is no \mathbf{Q} -homomorphism from $\mathcal{I}_{\mathcal{T}_2, t'} |_{\mathbf{Q}}^{\text{con}}$ to $\mathcal{I}_{L,1}$ that maps the root to a ρ -neighbor of n .

Note that this is sufficient since the subtrees of $\mathcal{I}_{\mathcal{T}_2, t}$ rooted at ρ -successors of the root are exactly the models $\mathcal{I}_{\mathcal{T}_2, t'}$, $t \rightsquigarrow_{\rho}^{\mathcal{T}_2} t'$. Condition (ii) is recursive in the sense that we are faced with the same conditions that we started with, only for a different \mathcal{T}_2 -type and element of $\text{ind}(\mathcal{A})$. We need to ensure that the recursion terminates, which is achieved by assigning appropriate priorities to states.

The above is implemented as follows. Define $\mathfrak{A}_4 = (Q_4, \Theta, q_0, \delta_4, \Omega_4)$ where

$$Q_4 = \{q_0, q_1\} \cup \{\bar{q}_2^t, \bar{q}_3^t, \bar{q}_{3b}^t \mid t \in \text{tp}(\mathcal{T}_2)\} \cup \{\bar{q}_2^{\rho, t}, \bar{q}_2^{\rho, t, \downarrow} \mid t \in \text{tp}(\mathcal{T}_2), \rho \text{ set of } \text{sig}(\mathcal{T}_2)\text{-roles}\},$$

and Ω_4 assigns zero to all states, except for states of the form q_0 and q_2^t , $t \in \text{tp}(\mathcal{T}_2)$, to which it assigns one. For a \mathcal{T}_2 -type t and set of roles ρ , we use $t|_{\mathbf{Q}}$ and $\rho|_{\mathbf{Q}}$ to denote the restriction

of t and ρ to (the elements that only use symbols from) \mathbf{Q} . For each $\sigma = (L_0, L_1, L_2)$, δ_4 contains the following transitions:

$$\begin{aligned}
 \delta_4(q_0, \sigma) &= \begin{cases} \diamond q_0 \vee q_1 \vee \bar{q}_2^{L_2 \cap \mathbf{N}_C} \vee \bigvee_{t' \in R_{\mathbf{Q}}(L_2 \cap \mathbf{N}_C)} \bar{q}_3^{t'} & \text{if } L_0 \neq \emptyset \\ \text{false} & \text{otherwise} \end{cases} \\
 \delta_4(q_1, \sigma) &= \begin{cases} \text{true} & \text{if } L_2 \setminus L_1 \text{ contains a } \mathbf{Q}\text{-role} \\ \text{false} & \text{otherwise} \end{cases} \\
 \delta_4(\bar{q}_2^t, \sigma) &= \begin{cases} \text{true} & \text{if } t|_{\mathbf{Q}} \not\subseteq L_1 \\ \bigvee_{t \rightsquigarrow_{\rho} t'} \bar{q}_2^{\rho, t'} & \text{otherwise} \end{cases} \\
 \delta_4(\bar{q}_2^{\rho, t}, \sigma) &= \begin{cases} \Box \bar{q}_2^{\rho, t, \downarrow} & \text{if } \rho^-|_{\mathbf{Q}} \not\subseteq L_1 \\ \Box \bar{q}_2^{\rho, t, \downarrow} \wedge \diamond^- \bar{q}_2^t & \text{if } \rho^-|_{\mathbf{Q}} \subseteq L_1 \end{cases} \\
 \delta_4(\bar{q}_2^{\rho, t, \downarrow}, \sigma) &= \begin{cases} \text{true} & \text{if } \rho|_{\mathbf{Q}} \not\subseteq L_1 \\ \bar{q}_2^t & \text{if } \rho|_{\mathbf{Q}} \subseteq L_1 \end{cases} \\
 \delta_4(\bar{q}_3^t, \sigma) &= \Box \bar{q}_3^t \wedge \Box^- \bar{q}_3^t \wedge \bar{q}_2^t \wedge \bar{q}_{3b}^t \\
 \delta_4(\bar{q}_{3b}^t, \sigma) &= \begin{cases} \text{true} & \text{if } L_0 = \emptyset \text{ or } \mathcal{I}_{\mathcal{T}_2, t}|_{\mathbf{Q}}^{\text{con}} \not\stackrel{\text{fin}}{\dashv} \mathcal{I}_{\mathcal{T}_1, L_1 \cap \mathbf{N}_C} \\ \text{false} & \text{otherwise} \end{cases}
 \end{aligned}$$

Note that states with index $i \in \{1, 2, 3, 3b\}$ are used to enforce Condition i from above. Also note that bounded homomorphisms are not handled directly by the automaton. The condition $\mathcal{I}_{\mathcal{T}_2, t}|_{\mathbf{Q}}^{\text{con}} \not\stackrel{\text{fin}}{\dashv} \mathcal{I}_{\mathcal{T}_1, L_1 \cap \mathbf{N}_C}$ from the last transition is only needed during the construction of the automaton, and it can be decided in the required time due to Theorem 21. The sets $R_{\mathbf{Q}}(t)$ can clearly be computed in single exponential time.

AUTOMATON \mathfrak{A}_3 .

Recall that \mathfrak{A}_3 has to accept an input (T, L) iff \mathcal{A}_L is consistent with \mathcal{T}_2 , and $\mathcal{I}_{L, 2}$ is $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}_L}$ restricted to $\text{ind}(\mathcal{A}_L)$. The construction requires a few preliminaries, in particular a characterization of whether $\mathcal{T}, \mathcal{A} \models A(a)$ in terms of derivation trees. Similar yet slightly different characterizations have been used before (e.g. in Bienvenu, Lutz, & Wolter, 2013).

Let \mathcal{T} be a Horn- \mathcal{ALCHIF} TBox and \mathcal{A} an ABox. A *derivation tree* for an assertion $A_0(a_0)$ in \mathcal{A} w.r.t. \mathcal{T} with $A_0 \in \mathbf{N}_C$ is a finite $\text{ind}(\mathcal{A}) \times \mathbf{N}_C$ -labeled tree (T, V) that satisfies the following conditions:

1. $V(\varepsilon) = (a_0, A_0)$;
2. if $V(x) = (a, A)$ and neither $A(a) \in \mathcal{A}$ nor $\top \sqsubseteq A \in \mathcal{T}$, then one of the following holds:
 - (i) x has successors y_1, \dots, y_k , $k \geq 1$ with $V(y_i) = (a, A_i)$ for $1 \leq i \leq k$ and $\mathcal{T} \models A_1 \sqcap \dots \sqcap A_k \sqsubseteq A$;
 - (ii) x has a single successor y with $V(y) = (b, B)$ and there is an $B \sqsubseteq \forall r. A \in \mathcal{T}$ and an $s(b, a) \in \mathcal{A}$ such that $\mathcal{T} \models s \sqsubseteq r$;
 - (iii) x has a single successor y with $V(y) = (b, B)$ and there is a $B \sqsubseteq \exists r. A \in \mathcal{T}$ such that $r(b, a) \in \mathcal{A}$ and $\text{func}(r) \in \mathcal{T}$.

Item (i) of Point 2 above requires $\mathcal{T} \models A_1 \sqcap \dots \sqcap A_n \sqsubseteq A$ instead of $A_1 \sqcap A_2 \sqsubseteq A \in \mathcal{T}$ to ‘shortcut’ parts of the universal model that are generated by existential restrictions. In fact, elements of the universal model generated by existential restrictions do never appear in a derivation tree. The main properties of derivation trees are summarized in the following lemma, proved in Appendix B.2.

Lemma 26 *Let \mathcal{T} be a Horn-ALCHIF TBox and \mathcal{A} . Then*

1. *if \mathcal{A} is consistent with \mathcal{T} , then $\mathcal{T}, \mathcal{A} \models A_0(a_0)$ iff there is a derivation tree for $A_0(a_0)$ in \mathcal{A} w.r.t. \mathcal{T} , for all assertions $A_0(a_0)$;*
2. *\mathcal{A} is consistent with \mathcal{T} iff the following are satisfied:*
 - (a) *the ABox $\mathcal{A}_a = \{A(a) \mid \mathcal{T}, \mathcal{A} \models A(a)\}$ is consistent with \mathcal{T} , for all $a \in \text{ind}(\mathcal{A})$;*
 - (b) *the relation $\{(a, b) \mid s(a, b) \in \mathcal{A}\}$ is a partial function whenever $\text{func}(s) \in \mathcal{T}$.*

We now construct the automaton \mathfrak{A}_3 . It ensures that when a Θ -labeled tree (T, L) is accepted, then for all $n \in \text{ind}(\mathcal{A}_L)$, concept names A , and roles r :

- (i) $A \in L_2(n)$ iff there is a derivation tree for $A(n)$ in \mathcal{A}_L ;
- (ii) for all $n \neq \varepsilon$, $r \in L_2(n)$ iff there is a role s such that $s(n \cdot -1, n) \in \mathcal{A}_L$ and $\mathcal{T}_2 \models s \sqsubseteq r$.

By Point 1 of Lemma 26, these conditions ensure that the interpretation $\mathcal{I}_{L,2}$ is the universal model of \mathcal{T}_2 and \mathcal{A} restricted to $\text{ind}(\mathcal{A}_L)$, in case \mathcal{A} is consistent with \mathcal{T}_2 . Given this, we can verify consistency of \mathcal{A} with \mathcal{T}_2 based on Point 2 of Lemma 26, that is, we ensure the following:

- (iii) the set $L_2(n) \cap \mathbf{N}_C$ is consistent with \mathcal{T}_2 , for all $n \in \text{ind}(\mathcal{A})$;
- (iv) the relation $\{(a, b) \mid s(a, b) \in \mathcal{A}_L\}$ is a partial function whenever $\text{func}(s) \in \mathcal{T}_2$.

We take $\mathfrak{A}_3 = (Q_3, \Theta, q_0, \delta_3, \Omega_3)$ where

$$\begin{aligned} Q_3 = & \{q_0, q'_0, q_1\} \cup \{q_A, \bar{q}_A \mid A \in \Theta_2 \cap \mathbf{N}_C\} \cup \\ & \{q_r, \bar{q}_r, q_r^A, \bar{q}_r^A, q_r^f, q_{\neg r} \mid r \in \Theta_2 \setminus \mathbf{N}_C\} \cup \\ & \{q_{r,B}, \bar{q}_{r,B} \mid r \in \Theta_2 \cap \mathbf{N}_R, B \in \Theta_2 \cap \mathbf{N}_C\} \end{aligned}$$

and Ω_3 assigns zero to all states, except for states of the form q_A , to which it assigns 1.

For Condition (i), we use states q_A for the “ \Leftarrow ” part, and states \bar{q}_A for the “ \Rightarrow ” part. Intuitively, a state q_A assigned to some node n is an obligation to verify the existence of a derivation tree for $A(n)$. Conversely, \bar{q}_A is the obligation that there is *no* such derivation tree. Conditions (ii) to (iv) are rather straightforward to verify. The automaton starts with

the following transitions for every $\sigma = (L_0, L_1, L_2)$:

$$\begin{aligned}
 \delta_3(q_0, \sigma) &= \begin{cases} \text{true} & \text{if } L_0 = \emptyset \\ q'_0 & \text{if } L_0 \neq \emptyset \end{cases} \\
 \delta_3(q'_0, \sigma) &= \begin{cases} \text{false} & \text{if } L_2 \cap \mathbf{N}_C \text{ inconsistent with } \mathcal{T}_2 \\ \square q_0 \wedge \square q_1 \wedge \bigwedge_{A \in L_2 \cap \mathbf{N}_C} q_A \wedge \bigwedge_{A \in (\Theta_2 \cap \mathbf{N}_C) \setminus L_2} \bar{q}_A & \text{otherwise} \end{cases} \\
 \delta_3(q_1, \sigma) &= \begin{cases} \text{true} & \text{if } L_0 = \emptyset \\ \bigwedge_{\text{func}(r) \in \mathcal{T}_2} q_r^f \wedge \bigwedge_{r \in L_2 \cap \mathbf{N}_R} q_r \wedge \bigwedge_{r \in (\Theta_2 \cap \mathbf{N}_R) \setminus L_2} \bar{q}_r & \text{otherwise} \end{cases} \\
 \delta_3(q_r^f, \sigma) &= \begin{cases} \square q_{\neg r} & \text{if } r^- \in L_0 \\ \square_1 q_{\neg r} & \text{if } r^- \notin L_0 \end{cases} \\
 \delta_3(q_{\neg r}, \sigma) &= \begin{cases} \text{true} & \text{if } r \notin L_0 \\ \text{false} & \text{otherwise} \end{cases} \\
 \delta_3(q_r, \sigma) &= \begin{cases} \text{true} & \text{if there is an } s \notin L_0 \text{ with } \mathcal{T}_2 \models s \sqsubseteq r \\ \text{false} & \text{otherwise} \end{cases} \\
 \delta_3(\bar{q}_r, \sigma) &= \begin{cases} \text{true} & \text{if there is no } s \notin L_0 \text{ with } \mathcal{T}_2 \models s \sqsubseteq r \\ \text{false} & \text{otherwise} \end{cases}
 \end{aligned}$$

For states q_A , we implement the conditions of derivation trees as transitions. Finiteness of the derivation tree is ensured by the priority assigned to these states. The relevant transitions are as follows:

$$\begin{aligned}
 \delta_3(q_A, \sigma) &= \text{false} && \text{if } L_0 = \emptyset \\
 \delta_3(q_A, \sigma) &= \text{true} && \text{if } A \in L_0 \\
 \delta_3(q_A, \sigma) &= \bigvee_{\mathcal{T}_2 \models A_1 \sqcap \dots \sqcap A_n \sqsubseteq A} (q_{A_1} \wedge \dots \wedge q_{A_n}) \vee && \text{if } L_0 \neq \emptyset \text{ and } A \notin L_0 \\
 &\quad \bigvee_{B \sqsubseteq \forall r. A \in \mathcal{T}_2, \mathcal{T}_2 \models s \sqsubseteq r} ((q_s^A \wedge \diamond^- q_B) \vee \diamond q_{s^-, B}) \vee \\
 &\quad \bigvee_{B \sqsubseteq \exists r. A \in \mathcal{T}_2, \text{func}(r) \in \mathcal{T}_2} ((q_s^A \wedge \diamond^- q_B) \vee \diamond q_{s^-, B}) \\
 \delta_3(q_r^A, \sigma) &= \begin{cases} \text{true} & \text{if } r \in L_0 \\ \text{false} & \text{otherwise} \end{cases} \\
 \delta_3(q_{s,B}, \sigma) &= q_s^A \wedge q_B
 \end{aligned}$$

The transitions for \bar{q}_A are obtained by dualizing the ones for q_A . More precisely, for every q of the form q_A , q_r^A , and $q_{s,B}$, we define $\delta_3(\bar{q}, \sigma) = \bar{\delta}_3(q, \sigma)$, where $\bar{\varphi}$ is obtained from φ by exchanging \wedge with \vee , \diamond with \square , \diamond^- with \square^- , and true with false , and replacing every state p with \bar{p} .

5. Tree-Shaped CQs and Deductive Conservative Extensions

We consider the 1tCQ versions of entailment, inseparability, and conservative extensions in Horn- \mathcal{ALCHIF} (and fragments) as well as their deductive companions in \mathcal{ELHIF}_\perp (and fragments). The results are summed up by the following theorem.

Theorem 27 *The following problems are 2EXPTIME-complete:*

1. In *Horn-ALCHIF* and any of its fragments that contains \mathcal{ELI} or *Horn-ALC*: (\mathbf{A}, \mathbf{Q}) -1tCQ entailment, (\mathbf{A}, \mathbf{Q}) -1tCQ inseparability, and (\mathbf{A}, \mathbf{Q}) -1tCQ conservative extensions; this holds even when $\mathbf{A} = \mathbf{Q}$;
2. In \mathcal{ELHIF}_\perp and any of its fragments that contains \mathcal{ELI} : \mathbf{S} -deductive entailment, \mathbf{S} -deductive inseparability, and \mathbf{S} -deductive conservative extensions.

Moreover, all these problems can be solved in time $2^{2^{p(|\mathcal{T}_2| \log |\mathcal{T}_1|)}}$, where p is a polynomial.

It is worth remembering that in \mathcal{EL} , which is the fragment of \mathcal{ELI} that does not admit inverse roles, the reasoning problems mentioned in Point 2 of Theorem 27 are EXPTIME-complete (Lutz & Wolter, 2010). We find it remarkable that adding inverse roles causes a jump by one exponential. This has so far only been observed for reasoning problems that involve (non-tree shaped) conjunctive queries (e.g. in Bienvenu et al., 2016), but this is not the case for the problems in Point 2.

For Point 1, the lower bounds have been established by Botoeva et al. (2019). The upper bound is proved almost exactly as in Section 4.2. By Theorem 19, the only difference is that the automaton \mathfrak{A}_4 now needs to accept a Θ -labeled tree (T, L) iff $\mathcal{I}_{L,2} \not\leq_{\mathbf{S}} \mathcal{I}_{L,1}$. To achieve this, \mathfrak{A}_4 ensures that for some $n \in \text{ind}(\mathcal{A}_L)$, one of the following conditions hold:

1. there is a \mathbf{Q} -role r and an $n' \in \text{ind}(\mathcal{A}_L)$ such that $(n, n') \in r^{\mathcal{I}_{L,2}}$, but $(n, n') \notin r^{\mathcal{I}_{L,1}}$;
- 2.' $\mathcal{I}_{\mathcal{T}_2, L_2(n) \cap \mathbf{N}_c} \not\leq_{\mathbf{Q}} \mathcal{I}_{L,1}$ via a \mathbf{Q} -simulation that relates the root of $\mathcal{I}_{\mathcal{T}_2, L_2(n) \cap \mathbf{N}_c}$ to n .

This can be achieved by modifying the construction of \mathfrak{A}_4 from Section 4.2 as follows: drop all states \bar{q}_3^t , \bar{q}_{3b}^t , and $\bar{q}_2^{\rho,t}, \bar{q}_2^{\rho,t,\downarrow}$ with $|\rho| > 1$ (and all according transitions) and replace the transitions for q_0 and \bar{q}_2^t as follows:

$$\delta_4(q_0, \sigma) = \begin{cases} \diamond q_0 \vee q_1 \vee \bar{q}_2^{L_2 \cap \mathbf{N}_c} & \text{if } L_0 \neq \emptyset \\ \text{false} & \text{otherwise} \end{cases}$$

$$\delta_4(\bar{q}_2^t, \sigma) = \begin{cases} \text{true} & \text{if } t|_{\mathbf{Q}} \not\subseteq L_1 \\ \bigvee_{t \rightsquigarrow_{\rho} \mathcal{T}_2 t'} \bigvee_{r \in \rho} \bar{q}_2^{\{r\}, t'} & \text{otherwise.} \end{cases}$$

It thus remains to prove Point 2 of Theorem 27. The upper bound is a consequence of Point 1, Proposition 8, and the fact that 1tCQ evaluation in \mathcal{ELHIF}_\perp can be done in exponential time (Eiter et al., 2008). We turn to the lower bounds.

Theorem 28 *In \mathcal{ELI} , deciding (\mathbf{S}, \mathbf{S}) -1tCQ conservative extensions, \mathbf{S} -deductive conservative extensions, deductive \mathbf{S} -entailment, and deductive \mathbf{S} -inseparability are 2EXPTIME-hard. This even holds when $\mathbf{S} = \text{sig}(\mathcal{T}_1)$.*

Proof. We start with 1tCQ conservative extensions (not assuming $\mathbf{S} = \text{sig}(\mathcal{T}_1)$) and reduce from the following problem which has been shown to be 2EXPTIME-hard (Gutiérrez-Basulto et al., 2018, implicit in Lemma 8).

Input: \mathcal{ELI} TBox \mathcal{T} , signature \mathbf{S} , concept names A, B

Question: $\mathcal{I}_{\mathcal{T}, \{A(a)\}} \rightarrow_{\mathbf{S}} \mathcal{I}_{\mathcal{T}, \{B(a)\}}$?

For the reduction, let $\mathbf{S}, \mathcal{T}, A, B$ be an input. Define two TBoxes \mathcal{T}_1 and \mathcal{T}_2 and a signature \mathbf{S}' by taking

$$\begin{aligned}\mathcal{T}_1 &= \mathcal{T} \cup \{Y \sqsubseteq \exists s.B\}, \\ \mathcal{T}_2 &= \mathcal{T}_1 \cup \{Y \sqsubseteq \exists s.A\}, \\ \mathbf{S}' &= \mathbf{S} \cup \{Y, s\},\end{aligned}$$

for some fresh concept name Y and a fresh role name s . By Theorem 19, it suffices to establish the following claim.

Claim. $\mathcal{I}_{\mathcal{T},\{A(a)\}} \rightarrow_{\mathbf{S}} \mathcal{I}_{\mathcal{T},\{B(a)\}}$ iff $\mathcal{I}_{\mathcal{T}_2,\mathcal{A}} \preceq_{\mathbf{S}'} \mathcal{I}_{\mathcal{T}_1,\mathcal{A}}$ for all tree-shaped \mathbf{S}' -ABoxes \mathcal{A} .

“ \Rightarrow ”. Assume $\mathcal{I}_{\mathcal{T},\{A(a)\}} \rightarrow_{\mathbf{S}} \mathcal{I}_{\mathcal{T},\{B(a)\}}$ and let \mathcal{A} be a tree-shaped \mathbf{S}' -ABox. By construction of \mathcal{T}_1 , the universal model $\mathcal{I}_{\mathcal{T}_1,\mathcal{A}}$ can be obtained from $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ by adding to every c with $Y(c) \in \mathcal{A}$ an s -successor to a fresh isomorphic copy of $\mathcal{I}_{\mathcal{T},\{B(a)\}}$. Similarly, the universal model $\mathcal{I}_{\mathcal{T}_2,\mathcal{A}}$ can be obtained from $\mathcal{I}_{\mathcal{T},\mathcal{A}}$ by adding to every c with $Y(c) \in \mathcal{A}$ an s -successor to a copy of $\mathcal{I}_{\mathcal{T},\{B(a)\}}$ and another s -successor to a copy of $\mathcal{I}_{\mathcal{T},\{A(a)\}}$. The required simulation can be easily constructed based on the existing homomorphism $\mathcal{I}_{\mathcal{T},\{A(a)\}} \rightarrow_{\mathbf{S}} \mathcal{I}_{\mathcal{T},\{B(a)\}}$.

“ \Leftarrow ”. Suppose that $\mathcal{I}_{\mathcal{T}_2,\mathcal{A}} \preceq_{\mathbf{S}'} \mathcal{I}_{\mathcal{T}_1,\mathcal{A}}$ for all tree-shaped \mathbf{S}' -ABoxes \mathcal{A} . Then, in particular, $\mathcal{I}_{\mathcal{T}_2,\{Y(a)\}} \preceq_{\mathbf{S}'} \mathcal{I}_{\mathcal{T}_1,\{Y(a)\}}$. We can now read off an \mathbf{S} -homomorphism from the \mathbf{S}' -simulation because $\mathcal{I}_{\mathcal{T}_2,\{Y(a)\}}$ and $\mathcal{I}_{\mathcal{T}_1,\{Y(a)\}}$ are tree-shaped.

This finishes the proof of the claim and shows that 1tCQ conservative extensions are 2EXPTIME-hard in \mathcal{ELI} . It remains to note that Proposition 9 provides a reduction from 1tCQ entailment over \mathcal{ELI} TBoxes to deductive conservative extensions in \mathcal{ELHIF}_\perp (importantly, we trivially have $\mathcal{T}_1 \models_{\mathbf{S}}^{\perp} \mathcal{T}_2$, no role hierarchies, and no functionality assertions).

For the case of $\mathbf{S} = \text{sig}(\mathcal{T}_1)$, we do a closer inspection of the 2EXPTIME-hardness proof of the mentioned homomorphism problem by Gutiérrez-Basulto et al. (2018). The constructed TBox \mathcal{T} is in fact a union of two disjoint TBoxes $\mathcal{T} = \mathcal{T}_1 \cup \mathcal{T}_2$ which additionally satisfy:

$$(*) \quad \mathcal{I}_{\mathcal{T}_1 \cup \mathcal{T}_2, \{A(a)\}} = \mathcal{I}_{\mathcal{T}_2, \{A(a)\}} \text{ for all } A \notin \text{sig}(\mathcal{T}_1), \text{ and } \mathcal{I}_{\mathcal{T}_1 \cup \mathcal{T}_2, \mathcal{A}} = \mathcal{I}_{\mathcal{T}_1, \mathcal{A}}, \text{ for all sig}(\mathcal{T}_1)\text{-ABoxes } \mathcal{A},$$

and the used signature is actually $\text{sig}(\mathcal{T}_1)$. Thus, the problem of deciding $\mathcal{I}_{\mathcal{T}_2, \{A(a)\}} \rightarrow_{\text{sig}(\mathcal{T}_1)} \mathcal{I}_{\mathcal{T}_1, \{B(a)\}}$, for $B \in \text{sig}(\mathcal{T}_1)$ and $A \notin \text{sig}(\mathcal{T}_1)$ is 2EXPTIME-hard for TBoxes $\mathcal{T}_1, \mathcal{T}_2$ satisfying (*). For the reduction, let us fix TBoxes \mathcal{T}'_1 and \mathcal{T}'_2 by taking

$$\begin{aligned}\mathcal{T}'_1 &= \mathcal{T}_1 \cup \{Y \sqsubseteq \exists s.B\}, \\ \mathcal{T}'_2 &= \mathcal{T}_1 \cup \mathcal{T}_2 \cup \{Y \sqsubseteq \exists s.B, Y \sqsubseteq \exists s.A\}, \\ \mathbf{S} &= \text{sig}(\mathcal{T}'_1) = \text{sig}(\mathcal{T}_1) \cup \{Y, s\}\end{aligned}$$

for fresh symbols Y, s . Based on (*), it is not hard to verify that the following claim can be proved analogously to the claim above, thus establishing hardness by Theorem 19.

Claim. $\mathcal{I}_{\mathcal{T}_2, \{A(a)\}} \rightarrow_{\text{sig}(\mathcal{T}_1)} \mathcal{I}_{\mathcal{T}_1, \{B(a)\}}$ iff $\mathcal{I}_{\mathcal{T}'_2, \mathcal{A}} \preceq_{\mathbf{S}} \mathcal{I}_{\mathcal{T}'_1, \mathcal{A}}$, for all tree-shaped \mathbf{S} -ABoxes. \square

6. Conclusion

We have shown that query conservative extensions of TBoxes are 2EXPTIME-complete for all DLs between \mathcal{ELI} and Horn- \mathcal{ALCHIF} as well as between Horn- \mathcal{ALC} and Horn- \mathcal{ALCHIF} , and that deductive conservative extensions are 2EXPTIME-complete for all DLs between \mathcal{ELI} and \mathcal{ELHIF}_\perp . This gives a fairly complete picture of the complexity and decidability of conservative extensions in Horn DLs with inverse roles.

An interesting problem left open is the decidability and complexity of deductive conservative extensions in Horn- \mathcal{ALCHIF} . A natural approach would be to come up with a characterization of the expressive power in terms of a ‘(bi)simulation-type’ indistinguishability relation on models, to then use that as the foundation for a characterization of deductive conservative extensions similar to what has been done for expressive DLs (Lutz & Wolter, 2011) and \mathcal{EL} (Lutz, Seylan, & Wolter, 2012), and to finally develop a decision procedure based on tree automata. A first step has recently been made by Jung, Papacchini, Wolter, and Zakharyashev (2019) who presented a (bi)simulation-type relation for the fragment Horn- \mathcal{ALC} of Horn- \mathcal{ALCHIF} . Unfortunately, the proposed relation is much more complicated than standard (bi)simulations and while it can be used to characterize deductive conservative extensions, it is far from clear how such a characterization can lead to a decision procedure.

It would also be interesting to add transitive roles to the picture, that is, to transition from Horn- \mathcal{ALCHIF} to Horn- \mathcal{SHIF} . In fact, we are not aware of any results on conservative extensions or related notions that concern description logics with transitive roles.

Finally, it would be interesting to consider both query and deductive conservative extensions in the Datalog $^\pm$ family of ontology languages (aka existential rules) such as frontier-guarded existential rules (Baget, Mugnier, Rudolph, & Thomazo, 2011). In fact, we are not aware of any (un)decidability results regarding conservative extensions in such languages. The increased existential power of Datalog $^\pm$ languages brings in serious additional technical challenges. In this context, it is interesting to remark that deductive conservative extensions are undecidable in the guarded fragment, which might be seen as an extension of relevant Datalog $^\pm$ languages (Jung, Lutz, Martel, Schneider, & Wolter, 2017).

Acknowledgments

This work was partially supported by DFG project SCHN 1234/3 and ERC Consolidator Grant 647289 CODA.

Appendix A. Proofs for Section 2

This section contains the proofs omitted from Sections 2.2 and 2.3.

A.1 Proof of Lemma 5

Lemma 5 *Let \mathcal{T}_1 and \mathcal{T}_2 be Horn- \mathcal{ALCHIF} TBoxes and let \mathbf{A}, \mathbf{Q} be signatures. Then \mathcal{T}_1 (\mathbf{A}, \mathbf{Q})-CQ entails \mathcal{T}_2 with inconsistent ABoxes iff one of the two following conditions holds.*

- (1) $\mathcal{T}_1 \models_{\mathbf{A}, \mathbf{Q}}^{\text{CQ}} \mathcal{T}_2$ and $\mathcal{T}_1 \models_{\mathbf{A}}^\perp \mathcal{T}_2$;

(2) \mathcal{T}_1 is (\mathbf{A}, \mathbf{Q}) -universal.

Proof. We prove both implications via contraposition.

“ \Rightarrow ”. Assume (1) and (2) are both false, i.e., \mathcal{T}_1 is not (\mathbf{A}, \mathbf{Q}) -universal and either (a) $\mathcal{T}_1 \not\models_{\mathbf{A}, \mathbf{Q}}^{\text{CQ}} \mathcal{T}_2$ or (b) $\mathcal{T}_1 \not\models_{\mathbf{A}}^{\perp} \mathcal{T}_2$. In case (a), \mathcal{T}_1 trivially does not (\mathbf{A}, \mathbf{Q}) -CQ entail \mathcal{T}_2 with inconsistent ABoxes. In case (b), consider a witness \mathbf{A} -ABox \mathcal{A} . Since \mathcal{T}_1 is not (\mathbf{A}, \mathbf{Q}) -universal, there is an \mathbf{A} -ABox \mathcal{A}' , a \mathbf{Q} -CQ $q(\mathbf{x})$ and a tuple $\mathbf{a} \subseteq \text{ind}(\mathcal{A}')$ with $|\mathbf{a}| = |\mathbf{x}|$ such that $\mathcal{T}_1, \mathcal{A}' \not\models q(\mathbf{a})$. We assume w.l.o.g. that \mathcal{A} and \mathcal{A}' use distinct sets of individuals. We set $\mathcal{A}'' = \mathcal{A} \cup \mathcal{A}'$ and have:

- $\mathcal{T}_2, \mathcal{A}'' \models q(\mathbf{a})$ because \mathcal{A} is inconsistent with \mathcal{T}_2 and so is \mathcal{A}'' ;
- $\mathcal{T}_1, \mathcal{A}'' \not\models q(\mathbf{a})$: let \mathcal{J} be the disjoint union of the universal model $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$ and the model \mathcal{I} witnessing $\mathcal{T}_1, \mathcal{A}' \not\models q(\mathbf{a})$. Clearly \mathcal{J} is a model of \mathcal{T}_1 and \mathcal{A}'' , but $\mathcal{J} \not\models q(\mathbf{a})$.

Hence \mathcal{T}_1 does not (\mathbf{A}, \mathbf{Q}) -CQ entail \mathcal{T}_2 with inconsistent ABoxes, as desired.

“ \Leftarrow ”. Assume \mathcal{T}_1 does not (\mathbf{A}, \mathbf{Q}) -CQ entail \mathcal{T}_2 with inconsistent ABoxes and consider a witness $(\mathcal{A}, q, \mathbf{a})$. Then it is immediate that (2) does not hold. Furthermore, if \mathcal{A} is consistent with both \mathcal{T}_1 and \mathcal{T}_2 , then $\mathcal{T}_1 \not\models_{\mathbf{A}, \mathbf{Q}}^{\text{CQ}} \mathcal{T}_2$. Otherwise \mathcal{A} must be inconsistent with \mathcal{T}_2 but consistent with \mathcal{T}_1 ; hence $\mathcal{T}_1 \not\models_{\mathbf{A}}^{\perp} \mathcal{T}_2$. Therefore (1) does not hold either. \square

A.2 Characterization of Inconsistency Entailment

We show that inconsistency entailment can be reduced to CQ entailment. We write $\mathcal{T}_1 \models_{\mathbf{A}}^{\text{fork}} \mathcal{T}_2$ if for all \mathbf{A} -ABoxes $\mathcal{A} = \{r(a, b), r(a, c)\}$: if \mathcal{A} is inconsistent with \mathcal{T}_2 , then also with \mathcal{T}_1 . Note that $\mathcal{T}_1 \models_{\mathbf{A}}^{\text{fork}} \mathcal{T}_2$ can be decided by evaluating the 1tCQ $A(a)$ on all \mathbf{A} -ABoxes of the form $\mathcal{A} = \{r(a, b), r(a, c)\}$ for both \mathcal{T}_i , where A is a concept name that does not occur in $\mathcal{T}_1 \cup \mathcal{T}_2$. Thus, $\mathcal{T}_1 \models_{\mathbf{A}}^{\text{fork}} \mathcal{T}_2$ is basically a shortcut for polynomially many CQ entailment tests.

Lemma 29 *Let \mathbf{A} be a signature and let \mathcal{T}_1 and \mathcal{T}_2 be Horn-ALCHIF TBoxes. Furthermore, let A be a fresh concept name and \mathcal{T}_i^A be obtained from \mathcal{T}_i by replacing each occurrence of \perp with A and adding the axioms $A \sqsubseteq \forall s.A$ and $A \sqsubseteq \forall s^-.A$ for every role s occurring in \mathcal{T}_i , for $i = 1, 2$. Then the following are equivalent.*

- (1) $\mathcal{T}_1 \models_{\mathbf{A}}^{\perp} \mathcal{T}_2$;
- (2) $\mathcal{T}_1^A \models_{\mathbf{A}, \{A\}}^{\text{CQ}} \mathcal{T}_2^A$ and $\mathcal{T}_1 \models_{\mathbf{A}}^{\text{fork}} \mathcal{T}_2$.

Proof. We prove both implications via contraposition.

(1) \Rightarrow (2). Assume $\mathcal{T}_1^A \not\models_{\mathbf{A}, \{A\}}^{\text{CQ}} \mathcal{T}_2^A$ or $\mathcal{T}_1 \not\models_{\mathbf{A}}^{\text{fork}} \mathcal{T}_2$. In case $\mathcal{T}_1 \not\models_{\mathbf{A}}^{\text{fork}} \mathcal{T}_2$, every witness ABox is a witness for $\mathcal{T}_1 \not\models_{\mathbf{A}}^{\perp} \mathcal{T}_2$ too.

In case $\mathcal{T}_1^A \not\models_{\mathbf{A}, \{A\}}^{\text{CQ}} \mathcal{T}_2^A$ is violated, consider a witness $(\mathcal{A}, q, \mathbf{a})$. Since A is the only symbol allowed in q , all atoms of q have the form $A(z)$ for arbitrary variables z . If q consists of several atoms, then it is disconnected and we can omit all but one atom from q and still have a witness (see also proof of Proposition 13, Property d). Hence we can assume w.l.o.g.

that q is of the form (i) $q(x) = A(x)$ or (ii) $q() = \exists y A(y)$ and, furthermore, that \mathcal{A} and thus the universal models $\mathcal{I}_{\mathcal{T}_i, \mathcal{A}}$ are connected. (Due to the “propagation” of A in the \mathcal{T}_i , we can even assume that q is of the form (i) only, but that does not matter in the following argumentation.) We now have:

- \mathcal{A} is inconsistent with \mathcal{T}_2 :

Assume to the contrary that \mathcal{A} is consistent with \mathcal{T}_2 and consider the universal model $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$ for \mathcal{T}_2 and \mathcal{A} (Section 2.5). Clearly, for all domain elements d of $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$, we have $\mathcal{T}_2 \not\models \prod \text{tp}_{\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}}(d) \sqsubseteq \perp$. Since A is fresh and by the definition of \mathcal{T}_2^A we get $\mathcal{T}_2^A \not\models \prod \text{tp}_{\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}}(d) \sqsubseteq A$. Now Lemma 11 (1) for \mathcal{T}_2^A implies that $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$ is a model of \mathcal{T}_2 and \mathcal{A} ; hence $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$ satisfies all axioms in \mathcal{T}_2^A that have been taken over from \mathcal{T}_2 without modification, i.e., all axioms that are not of the form $B \sqsubseteq A$. But axioms of the latter form are also satisfied because $\mathcal{T}_2^A \not\models \prod \text{tp}_{\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}}(d) \sqsubseteq A$ for every domain element d . Hence $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$ is a model of \mathcal{T}_2^A and \mathcal{A} . Now, since $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$ has no A -instance, we cannot have $\mathcal{T}_2^A, \mathcal{A} \models q(\mathbf{a})$ for any $\{A\}$ -query q ; contradicting the assumption that $(\mathcal{A}, q, \mathbf{a})$ is a witness.

- \mathcal{A} is consistent with \mathcal{T}_1 :

Since $(\mathcal{A}, q, \mathbf{a})$ is a witness, we have $\mathcal{I}_{\mathcal{T}_1^A, \mathcal{A}} \not\models q(\mathbf{a})$ by Lemma 11 (3). Due to the additional axioms in the definition of \mathcal{T}_1^A , which “propagate” A into every domain element of the connected (see above) universal model $\mathcal{I}_{\mathcal{T}_1^A, \mathcal{A}}$, we have $\mathcal{T}_1^A \not\models \prod \text{tp}_{\mathcal{I}_{\mathcal{T}_1^A, \mathcal{A}}}(d) \sqsubseteq A$ for all domain elements d . Since A is fresh, we have $\mathcal{T}_1 \not\models \prod \text{tp}_{\mathcal{I}_{\mathcal{T}_1^A, \mathcal{A}}}(d) \sqsubseteq \perp$. With the same reasoning as above, we obtain that $\mathcal{I}_{\mathcal{T}_1^A, \mathcal{A}}$ is a model of \mathcal{T}_1 and \mathcal{A} ; hence \mathcal{A} is consistent with \mathcal{T}_1 .

Consequently $\mathcal{T}_1 \not\models_{\mathbf{A}}^{\perp} \mathcal{T}_2$, as desired.

(2) \Rightarrow (1). Assume $\mathcal{T}_1 \not\models_{\mathbf{A}}^{\perp} \mathcal{T}_2$, i.e., there is an \mathbf{A} -ABox \mathcal{A} that is inconsistent with \mathcal{T}_2 but consistent with \mathcal{T}_1 . We need to show that $\mathcal{T}_1^A \not\models_{\mathbf{A}, \{A\}}^{\text{CQ}} \mathcal{T}_2^A$ or $\mathcal{T}_1 \not\models_{\mathbf{A}}^{\text{fork}} \mathcal{T}_2$.

From \mathcal{A} being inconsistent with \mathcal{T}_2 , we first conclude that one of the following two properties must hold.

- There is some $d \in B^{\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}}$ with $B \sqsubseteq \perp \in \mathcal{T}_2$;
- \mathcal{A} contains a “fork” $\mathcal{A}^- = \{r(a, b), r(a, c)\}$ such that \mathcal{A}^- is inconsistent with \mathcal{T}_2 .

Indeed, if neither (i) nor (ii) holds, then we have $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}} \models (\mathcal{T}_2, \mathcal{A})$, contradicting the inconsistency of \mathcal{A} with \mathcal{T}_2 : First, $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}} \models \mathcal{A}$ follows directly from the construction of $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$. Second, $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}} \models \mathcal{T}_2$ can be shown analogously to the (omitted) standard proof of Lemma 11 (1), via a case distinction over the axioms in \mathcal{T}_2 , using “not (i)” and “not (ii)” instead of the assumption that \mathcal{A} is consistent with \mathcal{T}_2 .

Now first assume that (ii) holds. Since \mathcal{A} is consistent with \mathcal{T}_1 , so is \mathcal{A}^- . Hence $\mathcal{T}_1 \not\models_{\mathbf{A}}^{\text{fork}} \mathcal{T}_2$.

In case (ii) does not hold, (i) must hold. To show that $\mathcal{T}_1^A \not\models_{\mathbf{A}, \{A\}}^{\text{CQ}} \mathcal{T}_2^A$, consider the CQ $q = A(x)$ and some $a \in \text{ind}(\mathcal{A})$ to which the element d from (i) is connected in $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$, i.e., if

$d \in \text{ind}(\mathcal{A})$, then choose $a = d$; otherwise choose a such that d is in the subtree $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}|_a$. We then have:

- \mathcal{A} is consistent with \mathcal{T}_2^A :

Since \mathcal{T}_2^A does not contain \perp and \mathcal{A} does not contain forks as in (ii), \mathcal{A} is consistent with \mathcal{T}_2^A is consistent, as witnessed by the universal model $\mathcal{I}_{\mathcal{T}_2^A, \mathcal{A}}$ (we again refer to the standard proof of Lemma 11 (1); except that the FA case in the ABox part of $\mathcal{I}_{\mathcal{T}_2^A, \mathcal{A}}$ is now due to “not (ii)”).

- \mathcal{A} is consistent with \mathcal{T}_1^A :

It is not difficult to see that $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$ is a model of \mathcal{T}_1^A and \mathcal{A} : by Lemma 11 (1), $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$ is a model of \mathcal{T}_1 and \mathcal{A} ; in particular, $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$ satisfies all axioms in \mathcal{T}_1^A that \mathcal{T}_1^A shares with \mathcal{T}_1 . The modified axioms $B \sqsubseteq A$ with $B \sqsubseteq \perp \in \mathcal{T}_1$ are satisfied, too, because $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$ cannot have any B -instances. Finally, the additional propagation axioms are satisfied because $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$ has no A -instance as A is fresh.

- $\mathcal{T}_2^A, \mathcal{A} \models q(a)$:

Due to (i), we have $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}} \models \exists y B(y)$ for some $B \sqsubseteq \perp \in \mathcal{T}_2$. Hence $\mathcal{I}_{\mathcal{T}_2^A, \mathcal{A}} \models \exists y B(y)$, which follows from the construction of both universal models (in fact the only difference between $\mathcal{I}_{\mathcal{T}_2^A, \mathcal{A}}$ and $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}$ is that some domain elements of $\mathcal{I}_{\mathcal{T}_2^A, \mathcal{A}}$ may be A -instances). Hence $\mathcal{I}_{\mathcal{T}_2^A, \mathcal{A}}$ has a B -instance in the subtree $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}}|_a$ and thus, by construction, an A -instance. By the “propagation” of A in \mathcal{T}_2^A , we have that a is an instance of A in $\mathcal{I}_{\mathcal{T}_2^A, \mathcal{A}}$; hence $\mathcal{I}_{\mathcal{T}_2^A, \mathcal{A}} \models A(a) = q$.

- $\mathcal{T}_1^A, \mathcal{A} \not\models q(a)$:

Follows from $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$ being a model of \mathcal{T}_1^A and \mathcal{A} (as shown above) and $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}} \not\models q(a)$ (given the lack of A -instances). □

A.3 Proof of Lemma 9

Lemma 9 *Let \mathbf{S} be a signature and $\mathcal{T}_1, \mathcal{T}_2 \in \mathcal{ELI}\mathcal{F}_\perp$ TBoxes such that $\mathcal{T}_1 \models_{\mathbf{S}, \mathbf{S}}^{\text{RI}} \mathcal{T}_2$. Then*

$$\mathcal{T}_1 \models_{\mathbf{S}}^{\text{TBox}} \mathcal{T}_2 \quad \text{iff} \quad \mathcal{T}_1 \models_{\mathbf{S}, \mathbf{S}}^{\text{1tCQ}} \mathcal{T}_2 \quad \text{and} \quad \mathcal{T}_1 \models_{\mathbf{S}}^{\perp} \mathcal{T}_2.$$

Proof. We prove both implications via contraposition.

“ \Leftarrow ”. We assume that $\mathcal{T}_1 \not\models_{\mathbf{S}}^{\text{TBox}} \mathcal{T}_2$. In case this is witnessed by an \mathbf{S} -FA $\text{func}(r)$, we immediately get a witness \mathbf{S} -ABox = $\{r(a, b), r(a, c)\}$ for $\mathcal{T}_1 \not\models_{\mathbf{S}}^{\perp} \mathcal{T}_2$ and are done.

Otherwise, \mathcal{T}_1 contains all \mathbf{S} -FAs from \mathcal{T}_2 , and there is a witness \mathbf{S} -CI $C \sqsubseteq D$ (witness RIs are excluded by the assumption $\mathcal{T}_1 \models_{\mathbf{S}, \mathbf{S}}^{\text{RI}} \mathcal{T}_2$). Since \mathcal{ELI}_\perp concepts that contain \perp are equivalent to \perp , the left-hand side C cannot contain \perp (i.e., is an \mathcal{ELI} concept) and, if D does, then $C \sqsubseteq \perp$ is a witness. We show that such witnesses give rise to either a witness \mathcal{A}_C for $\mathcal{T}_1 \not\models_{\mathbf{S}}^{\perp} \mathcal{T}_2$ or a witness (\mathcal{A}_C, q_D, a) for $\mathcal{T}_1 \not\models_{\mathbf{S}, \mathbf{S}}^{\text{1tCQ}} \mathcal{T}_2$ with $q_D(x)$ a 1tCQ .

We first consider the case that there is a witness $C \sqsubseteq \perp$ with C an \mathcal{ELI} concept. We can construct from C in the obvious way a tree-shaped **S**-ABox \mathcal{A}_C and root a : \mathcal{A} reflects the tree structure of C ; however, to respect the **S**-FAs in \mathcal{T}_1 (and thus those in \mathcal{T}_2), we need to merge the subtrees of all nodes that are r -neighbors of the same node, whenever $\text{func}(r) \in \mathcal{T}_1$. Consider the universal model $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}_C}$ ⁶ and observe that $a \in C^{\mathcal{I}_{\mathcal{T}_2, \mathcal{A}_C}}$ from the construction of $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}_C}$. Since $\mathcal{T}_2 \models C \sqsubseteq \perp$, we have that $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}_C}$ is not a model of \mathcal{T}_2 . Hence, by the contrapositive of Lemma 11 (1), \mathcal{A}_C is inconsistent with \mathcal{T}_2 . On the other hand, since $\mathcal{T}_1 \not\models C \sqsubseteq \perp$, there is a model $\mathcal{I} \models \mathcal{T}_1$ and an instance $d \in C^{\mathcal{I}}$. We can turn \mathcal{I} into a model of \mathcal{A}_C by interpreting the ABox individuals accordingly (“partial” unraveling might be necessary to ensure that the standard name assumption is respected), witnessing the consistency of \mathcal{A}_C with \mathcal{T}_1 . We thus have $\mathcal{T}_1 \not\models_{\mathbf{S}}^{\perp} \mathcal{T}_2$ and are done.

In the second case, *all* witnesses $C \sqsubseteq D$ consist solely of \mathcal{ELI} concepts C, D . We construct the same ABox \mathcal{A}_C with root a from C and transform D into an **S**-1tCQ $q_D(x)$ with a single answer variable that represents the tree shape of D . Now (\mathcal{A}_C, q_D, a) is a witness to $\mathcal{T}_1 \not\models_{\mathbf{S}, \mathbf{S}}^{\text{1tCQ}} \mathcal{T}_2$ for the following reasons.

- \mathcal{A}_C is consistent with \mathcal{T}_1 : a model can be obtained in the obvious way from the model witnessing $\mathcal{T}_1 \not\models C \sqsubseteq D$ (possibly involving “partial” unraveling as above).
- \mathcal{A}_C is consistent with \mathcal{T}_2 : since $C \sqsubseteq \perp$ is not a witness to $\mathcal{T}_1 \not\models_{\mathbf{S}}^{\text{TBox}} \mathcal{T}_2$, there must be a model $\mathcal{I} \models \mathcal{T}_2$ with $d \in C^{\mathcal{I}}$. We claim that we can turn \mathcal{I} into a model of \mathcal{A}_C by interpreting the ABox individuals without violating the standard name assumption. If we assume to the contrary that this is not possible, then there are subconcepts C_1, \dots, C_n of C corresponding to subtrees that have been merged in the construction of \mathcal{A}_C , such that $\mathcal{T}_2 \models C_1 \sqcap \dots \sqcap C_n \sqsubseteq \perp$. However, $\mathcal{T}_1 \not\models C_1 \sqcap \dots \sqcap C_n \sqsubseteq \perp$ because \mathcal{A}_C is consistent with \mathcal{T}_1 , as shown previously. Hence $C_1 \sqcap \dots \sqcap C_n \sqsubseteq \perp$ would be a witness to $\mathcal{T}_1 \not\models_{\mathbf{S}}^{\text{TBox}} \mathcal{T}_2$, which we have ruled out – a contradiction.
- $\mathcal{T}_2, \mathcal{A}_C \models q_D(a)$, witnessed by $\mathcal{I}_{\mathcal{T}_2, \mathcal{A}_C}$, together with $a \in C^{\mathcal{I}_{\mathcal{T}_2, \mathcal{A}_C}}$ and $\mathcal{T}_2 \models C \sqsubseteq D$.
- $\mathcal{T}_1, \mathcal{A}_C \not\models q_D(a)$: take a model \mathcal{I} witnessing $\mathcal{T}_1 \not\models C \sqsubseteq D$ and an element $d \in C^{\mathcal{I}} \setminus D^{\mathcal{I}}$. As in the previous case, we can turn \mathcal{I} into a model \mathcal{J} of \mathcal{A}_C by interpreting the ABox individuals (again involving unraveling if necessary), obtaining $\mathcal{J} \not\models q_D(a)$.

“ \Rightarrow ”. Assume $\mathcal{T}_1 \not\models_{\mathbf{S}, \mathbf{S}}^{\text{1tCQ}} \mathcal{T}_2$ or $\mathcal{T}_1 \not\models_{\mathbf{S}}^{\perp} \mathcal{T}_2$.

In case $\mathcal{T}_1 \not\models_{\mathbf{S}}^{\perp} \mathcal{T}_2$, consider a witness **S**-Box \mathcal{A} and assume w.l.o.g. that \mathcal{A} is tree-shaped. Let $a \in \text{ind}(\mathcal{A})$ be its root. We can assume that \mathcal{T}_1 contains all **S**-FAs from \mathcal{T}_2 (otherwise $\mathcal{T}_1 \not\models_{\mathbf{S}}^{\text{TBox}} \mathcal{T}_2$ and we are done). We turn \mathcal{A} into an **S**- \mathcal{ELI} concept $C_{\mathcal{A}}$ in the obvious way. Then $C_{\mathcal{A}} \sqsubseteq \perp$ is a witness to $\mathcal{T}_1 \not\models_{\mathbf{S}}^{\text{TBox}} \mathcal{T}_2$:

- $\mathcal{T}_2 \models C_{\mathcal{A}} \sqsubseteq \perp$ because, if there were a model \mathcal{I} of \mathcal{T}_2 with $d \in C_{\mathcal{A}}^{\mathcal{I}}$, we could turn it into a model of \mathcal{T}_2 and \mathcal{A} by interpreting the ABox individuals accordingly (possibly involving partial unraveling as above), which would contradict the assumption that \mathcal{A} is a witness to $\mathcal{T}_1 \not\models_{\mathbf{S}}^{\perp} \mathcal{T}_2$.

6. The assumption that \mathcal{A} is consistent with \mathcal{T} is not needed for the construction of $\mathcal{I}_{\mathcal{T}, \mathcal{A}}$, only for the proof of Lemma 11 (1).

- $\mathcal{T}_1 \not\sqsubseteq_{\mathbf{S}} C_{\mathcal{A}} \sqsubseteq \perp$, witnessed by $\mathcal{I}_{\mathcal{T}_1, \mathcal{A}}$.

In case $\mathcal{T}_1 \not\sqsubseteq_{\mathbf{S}}^{\text{1tCQ}} \mathcal{T}_2$, by Proposition 14 there is a witness (\mathcal{A}, q, a) with \mathcal{A} tree-shaped and q an **S**-1tCQ. We construct $C_{\mathcal{A}}$ as above and another **S-ELI** concept D_q from q in the obvious way. It can be shown analogously to the previous case that $C_{\mathcal{A}} \sqsubseteq D_q$ is a witness to $\mathcal{T}_1 \not\sqsubseteq_{\mathbf{S}}^{\text{TBox}} \mathcal{T}_2$. \square

Appendix B. Proofs for Section 4

We first state the semantics of 2ATA_c precisely. Let (T, L) be a Θ -labeled tree and $\mathfrak{A} = (Q, \Theta, q_0, \delta, \Omega)$ a 2ATA_c . A *run of \mathfrak{A} over (T, L)* is a $T \times Q$ -labeled tree (T_r, r) such that $\varepsilon \in T_r$, $r(\varepsilon) = (\varepsilon, q_0)$, and for all $y \in T_r$ with $r(y) = (x, q)$ and $\delta(q, V(x)) = \theta$, there is an assignment v of truth values to the transition atoms in θ such that v satisfies θ and:

- if $v(q') = 1$, then $r(y') = (x, q')$ for some successor y' of y in T_r ;
- if $v(\diamond^- q') = 1$, then $x \neq \varepsilon$ and $r(y') = (x \cdot -1, q')$ for some successor y' of y in T_r ;
- if $v(\square^- q') = 1$, then $x = \varepsilon$ or $r(y') = (x \cdot -1, q')$ for some successor y' of y in T_r ;
- if $v(\diamond_n q') = 1$, then there are pairwise different i_1, \dots, i_n such that, for each j , there is some successor y' of y in T_r with $r(y') = (x \cdot i_j, q')$;
- if $v(\square_n q') = 1$, then for all but n successors x' of x , there is a successor y' of y in T_r with $r(y') = (x', q')$.

Let $\gamma = i_0 i_1 \dots$ be an infinite path in T_r and denote, for all $j \geq 0$, with q_j the state such that $r(i_j) = (x, q_j)$. The path γ is *accepting* if the largest number m such that $\Omega(q_j) = m$ for infinitely many j is even. A run (T_r, r) is accepting, if all infinite paths in T_r are accepting. \mathfrak{A} accepts a tree if \mathfrak{A} has an accepting run over it.

B.1 Proof of Theorem 24

The proof is by reduction to the emptiness problem of standard two-way alternating tree automata on trees of some fixed outdegree (Vardi, 1998). We need to introduce strategy trees similar to (Vardi, 1998, Section 4). A *strategy tree for \mathfrak{A}* is a tree (T, τ) where τ labels every node in T with a subset $\tau(x) \subseteq 2^{Q \times \mathbb{N} \cup \{-1\} \times Q}$, that is, with a graph with nodes from Q and edges labeled with natural numbers or -1 . Intuitively, $(q, i, p) \in \tau(x)$ expresses that, if we reached node x in state q , then we should send a copy of the automaton in state p to $x \cdot i$. For each label ζ , we define $\text{state}(\zeta) = \{q \mid (q, i, q') \in \zeta\}$, that is, the set of sources in the graph ζ . A strategy tree is *on an input tree (T', L)* if $T = T'$, $q_0 \in \text{state}(\tau(\varepsilon))$, and for every $x \in T$, the following conditions are satisfied:

- (i) if $(q, i, p) \in \tau(x)$, then $x \cdot i \in T$;
- (ii) if $(q, i, p) \in \tau(x)$, then $p \in \text{state}(\tau(x \cdot i))$;
- (iii) if $q \in \text{state}(\tau(x))$, then the truth assignment $v_{q,x}$ defined below satisfies $\delta(q, L(x))$:

- $v_{q,x}(p) = 1$ iff $(q, 0, p) \in \tau(x)$;
- $v_{q,x}(\diamond^- p) = 1$ iff $(q, -1, p) \in \tau(x)$;
- $v_{q,x}(\square^- p) = 1$ iff $x = \varepsilon$ or $(q, -1, p) \in \tau(x)$;
- $v_{q,x}(\diamond_n p) = 1$ iff $(q, i, p) \in \tau(x)$ for n pairwise distinct $i \geq 1$;
- $v_{q,x}(\square_n p) = 1$ iff for all but at most n values $i \geq 1$ with $x \cdot i \in T$, we have $(q, i, p) \in \tau(x)$.

A *path* β in a strategy tree (T, τ) is a sequence $\beta = (u_1, q_1)(u_2, q_2) \cdots$ of pairs from $T \times Q$ such that for all $i > 0$, there is some c_i such that $(q_i, c_i, q_{i+1}) \in \tau(u_i)$ and $u_{i+1} = u_i \cdot c_i$. Thus, β is obtained by moves prescribed in the strategy tree. We say that β is *accepting* if the largest number m such that $\Omega(q_i) = m$, for infinitely many i , is even. A strategy tree (T, τ) is *accepting* if all infinite paths in (T, τ) are accepting.

Lemma 30 *A \mathcal{LATA}_c accepts an input tree iff there is an accepting strategy tree on the input tree.*

Proof. The “ \Leftarrow ” direction is immediate: just read off an accepting run from the accepting strategy tree.

For the “ \Rightarrow ” direction, we observe that acceptance of an input tree can be defined in terms of a parity game between Player 1 (trying to show that the input is accepted) and Player 2 (trying to challenge that). The initial configuration is (ε, q_0) and Player 1 begins. Consider a configuration (x, q) . Player 1 chooses a satisfying truth assignment v of $\delta(q, L(x))$. Player 2 chooses an atom α with $v_{q,x}(\alpha) = 1$ and determines the next configuration as follows:

- if $\alpha = p$, then the next configuration is (x, p) ;
- if $\alpha = \diamond^- p$, then the next configuration is $(x \cdot -1, p)$ unless $x = \varepsilon$; in this case, Player 1 loses immediately;
- if $\alpha = \square^- p$, then the next configuration is $(x \cdot -1, p)$ unless $x = \varepsilon$; in this case, Player 2 loses immediately;
- if $\alpha = \diamond_n p$, then Player 1 selects pairwise distinct i_1, \dots, i_n with $x \cdot i_j \in T$, for all j (and loses if she cannot); Player 2 then chooses some i_j and the next configuration is $(x \cdot i_j, p)$;
- if $\alpha = \square_n p$, then Player 1 selects n values i_1, \dots, i_n ; Player 2 then chooses some $\ell \notin \{i_1, \dots, i_n\}$ such that $x \cdot \ell \in T$ (and loses if he cannot) and the next configuration is $(x \cdot \ell, p)$.

Player 1 wins an infinite play $(x_0, q_0)(x_1, q_1) \cdots$ if the largest number m such that $\Omega(q_i) = m$, for infinitely many i , is even. It is not difficult to see that Player 1 has a winning strategy on an input tree iff \mathfrak{A} accepts the input tree.

Observe now that the defined game is a parity game and thus Player 1 has a winning strategy iff she has a *memoryless* winning strategy (Emerson & Jutla, 1991). It remains to observe that a memoryless winning strategy is nothing else than an accepting strategy tree. \square

Lemma 31 *If $L(\mathfrak{A}) \neq \emptyset$, then there is some $(T, L) \in L(\mathfrak{A})$ such that T has outdegree at most $n \cdot C$, where n is the number of states in \mathfrak{A} and C is the largest number in (some transition \diamond_{mp} or \square_{mp} in) δ .*

Proof. Let $(T, L) \in L(\mathfrak{A})$ and τ an accepting strategy tree on T , and let C be the largest number appearing in δ . We inductively construct a tree (T', L') with $T' \subseteq T$ and L' the restriction of L to T' and an accepting strategy tree τ' on (T', L') . For the induction base, we start with $T' = \{\varepsilon\}$ and τ' the empty mapping. For the inductive step, assume that $\tau'(x)$ is still undefined for some $x \in T'$, and proceed as follows:

1. For every $(q, i, p) \in \tau(x)$ with $i \in \{-1, 0\}$, add $(q, i, p) \in \tau'(x)$;
2. for every $p \in Q$, define $N_p = \{i \geq 1 \mid (q, i, p) \in \tau(x), x \cdot i \in T\}$ and let $N'_p \subseteq N_p$ be a subset of N_p with precisely $\min(C, |N_p|)$ elements. Then:
 - (a) for all $i \in N'_p$, add $x \cdot i \in T'$;
 - (b) for all $(q, i, p) \in \tau(x)$ with $i \in N'_p$, add $(q, i, p) \in \tau'(x)$;
 - (c) for all $q \in \text{state}(x)$ and $i \in N'_p$, add $(q, i, p) \in \tau'(x)$.

By Step 2 above, T' has outdegree bounded by $|Q| \cdot C$. It remains to show that τ' is an accepting strategy tree on T' . Observe first that, by construction, $q_0 \in \text{state}(\tau'(\varepsilon))$.

We verify Conditions (i)–(iii) of a strategy tree being *on an input tree*. Condition (i) follows directly from the construction. For (ii), assume that $(q, i, p) \in \tau'(x)$. By construction, there is some q' with $(q', i, p) \in \tau(x)$, and, by Condition (ii) $p \in \text{state}(\tau(x \cdot i))$. Hence, there is some $(p, j, p') \in \text{state}(\tau(x \cdot i))$. By construction, there is also some $(p, j', p') \in \text{state}(\tau'(x \cdot i))$, thus $p \in \text{state}(x \cdot i)$. For Condition (iii), take any $x \in T'$ and $q \in \text{state}(\tau'(x))$. As $q \in \text{state}(\tau(x))$, we know that the truth assignment $v_{q,x}$ defined for τ in Condition (iii) satisfies $\delta(q, L(x))$. We show that for all transitions α with $v_{q,x}(\alpha) = 1$, we also have $v'_{q,x}(\alpha) = 1$, where $v'_{q,x}$ is the truth assignment defined for τ' . By Step 1 of the construction, this is true for all α of the shape p , \diamond^-p , and \square^-p . Let now be $\alpha = \diamond_k p$, that is, there are k pairwise distinct $i \geq 1$ such that $(q, i, p) \in \tau(x)$. By the choice of C , we have $|N'_p| \geq k$. By Step 2(c), we know that there are k pairwise distinct i such that $(q, i, p) \in \tau'(x)$, hence $v'_{q,x}(\alpha) = 1$. Consider now $\alpha = \square_k p$, that is, for all but at most k values $i \geq 1$ with $x \cdot i \in T$, we have $(q, i, p) \in \tau(x)$. By Step 2(b), this remains true for τ' , hence $v'_{q,x}(\alpha) = 1$.

We finally argue that τ' is also accepting. Let $\beta = (u_1, q_1)(u_2, q_2) \cdots$ be an infinite path in (T', τ') . We construct an infinite path $\beta' = (u'_1, q_1)(u'_2, q_2)(u'_3, q_3) \cdots$ in (T, τ) as follows:

- $u'_1 = u_1$;
- Let $u_{i+1} = u_i \cdot \ell$ for some ℓ with $(q_i, \ell, q_{i+1}) \in \tau'(x)$. If $\ell \in \{0, 1\}$, we have $(q_i, \ell, q_{i+1}) \in \tau(x)$, by Step 1. We set $u'_{i+1} = u'_i \cdot \ell$. If $\ell \geq 0$ then, by Step 2(c), there is some ℓ' with $(q_i, \ell', q_{i+1}) \in \tau(x)$ and $x \cdot \ell' \in T'$. Set $u'_{i+1} = u'_i \cdot \ell'$.

Since every infinite path in (T, τ) is accepting, so is β' , and thus β . □

We now reduce the emptiness problem of $2ATA_c$ to the emptiness of alternating automata running on trees of fixed outdegree (Vardi, 1998), which we recall here. A tree T is k -ary if every node has exactly k . A *two-way alternating tree automaton over k -ary trees* ($2ATA^k$) that are Θ -labeled is a tuple $\mathcal{A} = (Q, \Theta, q_0, \delta, \Omega)$ where Q is a finite set of *states*, Θ is the *input alphabet*, $q_0 \in Q$ is an *initial state*, δ is the *transition function*, and $\Omega : Q \rightarrow \mathbb{N}$ is a *priority function*. The transition function maps a state q and some input letter θ to a *transition condition* $\delta(q, \theta)$, which is a positive Boolean formula over the truth constants **true**, **false**, and transitions of the form $(i, q) \in [k] \times Q$ where $[k] = \{-1, 0, \dots, k\}$. A *run* of \mathcal{A} on a Θ -labeled tree (T, L) is a $T \times Q$ -labeled tree (T_r, r) such that

1. $r(\varepsilon) = (\varepsilon, q_0)$;
2. for all $x \in T_r$ with $r(x) = (w, q)$ and $\delta(q, \tau(w)) = \varphi$, there is a (possibly empty) set $\mathcal{S} = \{(m_1, q_1), \dots, (m_n, q_n)\} \subseteq [k] \times Q$ such that \mathcal{S} satisfies φ and for $1 \leq i \leq n$, we have $x \cdot i \in T_r$, $w \cdot m_i$ is defined, and $\tau_r(x \cdot i) = (w \cdot m_i, q_i)$.

Accepting runs and accepted trees are defined as for $2ATA_c$ s. It is well-known that emptiness for $2ATA^k$ s can be checked in exponential time (Vardi, 1998), more precisely:

Theorem 32 *The emptiness problem for $2ATA^k$ s can be solved in time single exponential in the number of states and the maximal priority, and polynomial in all other inputs.*

We are now ready to prove Theorem 24.

Theorem 24 *The emptiness problem for $2ATA_c$ is in EXPTIME. More precisely, it can be solved in time single exponential in the number of states and the maximal priority, and polynomial in all other components.*

Proof. Let $\mathfrak{A} = (Q, \Theta, q_0, \delta, \Omega)$ be an $2ATA_c$ with n states and C the largest number in δ . We devise a $2ATA^k$ $\mathfrak{A}' = (Q', \Theta', q'_0, \delta', \Omega)$ with $k = n \cdot C$, that is, the bound from Lemma 31, such that $L(\mathfrak{A})$ is empty iff $L(\mathfrak{A}')$ is empty. Set $Q' = Q \cup \{q'_0, q_1, q_r, q_\perp\}$ and $\Theta' = (\Theta \cup \{d_\perp\}) \times \{0, 1\}$. The extended alphabet and the extra states are used to simulate transitions of the form $\square^- p$ (using the second component in the alphabet and state q_r) and to allow for input trees of outdegree less than k (using the additional symbol d_\perp).

We obtain δ' from δ by replacing q with $(0, q)$, $\diamond^- q$ with $(-1, q)$ and $\square^- q$ with $(0, q_r) \vee (-1, q)$. Moreover, we replace

- $\diamond_n q$ with $\bigvee_{X \in \binom{\{1, \dots, k\}}{n}} \bigwedge_{i \in X} (i, q)$;
- $\square_n q$ with $\bigvee_{X \in \binom{\{1, \dots, k\}}{n}} \bigwedge_{i \in \{1, \dots, N\} \setminus X} (i, q)$,

where, as usual, $\binom{M}{m}$ denotes the set of all m -elementary subsets of a set M . The transition function δ' for \mathfrak{A}' is δ extended with the following transitions:

$$\begin{aligned} \delta'(q'_0, (\theta, b)) &= \begin{cases} \text{false} & \text{if } b = 0 \text{ or } \theta = d_\perp, \\ q_0 \wedge \bigwedge_{i=1}^k (i, q_1) & \text{otherwise,} \end{cases} \\ \delta'(q_1, (\theta, b)) &= \begin{cases} \bigwedge_{i=1}^k (i, q_1) & \text{if } b = 0 \text{ and } \theta \neq d_\perp, \\ \text{true} & \text{if } \theta = d_\perp, \\ \text{false} & \text{otherwise,} \end{cases} \\ \delta'(q_r, (\theta, b)) &= \begin{cases} \text{true} & \text{if } b = 1, \\ \text{false} & \text{otherwise,} \end{cases} \\ \delta(q, (d_\perp, b)) &= \text{false} \quad \text{for all } q \in Q. \end{aligned}$$

By Lemma 31, we have that $L(\mathfrak{A})$ is empty iff $L(\mathfrak{A}')$ is empty. Note that this is not a polynomial time reduction to emptiness of 2ATA^k s due to the translations for \diamond_n and \square_n . It should be clear, however, that these translations can be *represented* in polynomial size, e.g., in the way they are given above. It remains to observe that, for the algorithm from Theorem 32 to work in the desired time, it suffices to be able to perform the following task in polynomial time:

- Given a set $\mathcal{S} \subseteq [k] \times Q$, a state $q \in Q$, and $a \in \Theta$, check whether \mathcal{S} satisfies $\delta(q, a)$.

This is certainly possible for the mentioned representation, which concludes the proof of the theorem. \square

B.2 Proof of Lemma 26

Lemma 26 *Let \mathcal{T} be a Horn-ALCHIF TBox and \mathcal{A} . Then*

1. *if \mathcal{A} is consistent with \mathcal{T} , then $\mathcal{T}, \mathcal{A} \models A_0(a_0)$ iff there is a derivation tree for $A_0(a_0)$ in \mathcal{A} w.r.t. \mathcal{T} , for all assertions $A_0(a_0)$;*

2. *\mathcal{A} is consistent with \mathcal{T} iff the following are satisfied:*

- (a) *the ABox $\mathcal{A}_a = \{A(a) \mid \mathcal{T}, \mathcal{A} \models A(a)\}$ is consistent with \mathcal{T} , for all $a \in \text{ind}(\mathcal{A})$;*
- (b) *the relation $\{(a, b) \mid s(a, b) \in \mathcal{A}\}$ is a partial function whenever $\text{func}(s) \in \mathcal{T}$.*

Proof. We start with Point 1. The “ \Leftarrow ” direction is straightforward to prove by induction on the depth of the derivation tree. We thus concentrate on the “ \Rightarrow ” direction. Thus assume that $\mathcal{T}, \mathcal{A} \models A_0(a_0)$. We construct a sequence of ABoxes by ‘chasing’ \mathcal{A} with the TBox \mathcal{T} , that is, by exhaustively and fairly applying the following rules:

1. If $A_1(a), \dots, A_k(a) \in \mathcal{A}$ and $\mathcal{T} \models A_1 \sqcap \dots \sqcap A_k \sqsubseteq A$, then add $A(a)$ to \mathcal{A} ;
2. If $r(a, b), A(a) \in \mathcal{A}$ and $A \sqsubseteq \forall r. B \in \mathcal{T}$, then add $B(b)$ to \mathcal{A} ;
3. If $r(a, b), B(a) \in \mathcal{A}$ and $B \sqsubseteq \exists r. A, \text{func}(r) \in \mathcal{T}$, then add $A(b)$ to \mathcal{A} .

Let $\mathcal{A} = \mathcal{A}_0, \mathcal{A}_1, \dots, \mathcal{A}_k$ be the emerging (finite!) sequence of ABoxes. We establish the following central claim.

Claim. There is a model \mathcal{I} of \mathcal{A}_k and \mathcal{T} such that for all $a \in \text{ind}(\mathcal{A})$ and concept names A , $a \in A^{\mathcal{I}}$ implies $A(a) \in \mathcal{A}_k$.

To construct the claimed \mathcal{I} , we start with \mathcal{I}_0 defined as follows.

$$\begin{aligned} \Delta^{\mathcal{I}_0} &= \text{ind}(\mathcal{A}) \\ A^{\mathcal{I}_0} &= \{a \mid A(a) \in \mathcal{A}_k\} \\ r^{\mathcal{I}_0} &= \{(a, b) \mid s(a, b) \in \mathcal{A}, \mathcal{T} \models s \sqsubseteq r\} \end{aligned}$$

Denote with \mathcal{A}_a the set $\{A(a) \in \mathcal{A}_k \mid A \in \mathbf{N}_C\}$. Now, obtain \mathcal{I} from \mathcal{I}_0 by proceeding as follows for every $a \in \text{ind}(\mathcal{A})$, $A \sqsubseteq \exists r.B \in \mathcal{T}$, and $a \in A^{\mathcal{I}}$. Let \mathcal{I}' be the subtree rooted at a ρ -successor b of a in $\mathcal{I}_{\mathcal{T}, \mathcal{A}_a}$ with $b \in B^{\mathcal{I}}$ and $r \in \rho$. If $\text{func}(r) \notin \mathcal{T}$ or there is no a' with $r(a, a') \in \mathcal{A}_k$, then add \mathcal{I}' as a ρ -successor of a in \mathcal{I}_0 . Clearly, \mathcal{I}_0 and thus \mathcal{I} is a model of \mathcal{A}_k . Based on the construction of \mathcal{A}_k and \mathcal{I} , it can also be verified that \mathcal{I} is a model of \mathcal{T} . This finishes the proof of the claim.

We are now ready to finish the proof of Point 1. From $\mathcal{T}, \mathcal{A} \models A_0(a_0)$ and the claim, we obtain $A_0(a_0) \in \mathcal{A}_k$. Exploiting that the three rules used to construct $\mathcal{A}_0, \dots, \mathcal{A}_k$ are in one-to-one correspondence with Conditions (i) to (iii) from the definition of derivation trees, it is easy to show by induction on i that for all $i \in \{1, \dots, k\}$ and all $A(a) \in \mathcal{A}_i$, there is a derivation tree for $A(a)$ in \mathcal{A} w.r.t. \mathcal{T} . In particular, this gives the desired derivation tree for $A_0(a_0)$.

Now for Point 2. The “ \Rightarrow ” direction is immediate, so we concentrate on “ \Leftarrow ”. Thus assume that Conditions (a) and (b) are satisfied. By (a) and Lemma 11 (1), the universal model $\mathcal{I}_a = \mathcal{I}_{\mathcal{T}, \mathcal{A}_a}$ is a model of \mathcal{T} and \mathcal{A}_a . In particular, \mathcal{I}_a is weakly tree-shaped with root a and satisfies, for all concept names $B \in \mathbf{N}_C$:

$$(*) \quad a \in B^{\mathcal{I}_a} \text{ iff } \mathcal{T}, \mathcal{A}_a \models B(a).$$

We construct a model \mathcal{I} of \mathcal{A} and \mathcal{T} as follows. Start with \mathcal{I}_0 by taking

$$\begin{aligned} \Delta^{\mathcal{I}_0} &= \text{ind}(\mathcal{A}) \\ A^{\mathcal{I}_0} &= \{a \mid A(a) \text{ has a derivation tree in } \mathcal{A} \text{ w.r.t. } \mathcal{T}\} \\ r^{\mathcal{I}_0} &= \{(a, b) \mid s(a, b) \in \mathcal{A}, \mathcal{T} \models s \sqsubseteq r\} \end{aligned}$$

and then adding, for every $a \in \text{ind}(\mathcal{A})$ and ρ -successor b of a in \mathcal{I}_a such that for all roles $r \in \rho$ with $\text{func}(r) \in \mathcal{T}$ we have $r(a, a') \notin \mathcal{A}$ for all $a' \in \text{ind}(\mathcal{A})$, the subinterpretation of \mathcal{I}_a rooted at b as a ρ -successor of a .

Clearly, \mathcal{I}_0 and thus \mathcal{I} is a model of \mathcal{A} . Based on (*) and the assumptions, it is also straightforward to show that \mathcal{I} is a model \mathcal{T} . By construction, all elements from $\Delta^{\mathcal{I}} \setminus \text{ind}(\mathcal{A})$ already satisfy \mathcal{T} . For $a \in \text{ind}(\mathcal{A})$, we distinguish cases on the shape of concept inclusions:

- If $\top \sqsubseteq A \in \mathcal{T}$, then trivially $A(a)$ has a derivation tree in \mathcal{A} w.r.t. \mathcal{T} , hence $a \in A^{\mathcal{I}}$.

- Suppose $A \sqsubseteq \perp \in \mathcal{T}$ and $a \in A^{\mathcal{I}}$. By definition of \mathcal{I} , $A(a)$ has a derivation tree in \mathcal{A} w.r.t. \mathcal{T} , and thus $\mathcal{T}, \mathcal{A} \models A(a)$. But then the ABox \mathcal{A}_a is not consistent with \mathcal{T} , a contradiction.
- If $A_1 \sqcap A_2 \sqsubseteq B \in \mathcal{T}$ and $a \in (A_1 \sqcap A_2)^{\mathcal{I}}$, then by construction, $A_1(a)$ and $A_2(a)$ have derivation trees in \mathcal{A} w.r.t. \mathcal{T} . Thus also $B(a)$ has a derivation tree, and hence $a \in B^{\mathcal{I}}$.
- If $A \sqsubseteq \exists r.B \in \mathcal{T}$ and $a \in A^{\mathcal{I}}$, then $A(a)$ has a derivation tree in \mathcal{A} w.r.t. \mathcal{T} . Thus, $\mathcal{T}, \mathcal{A} \models A(a)$, and a has a ρ -successor b in \mathcal{I}_a with $r \in \rho$ and $b \in B^{\mathcal{I}}$. By construction, a has the ρ -successor b in \mathcal{I} unless $\text{func}(r) \in \mathcal{T}$ and $r(a, a') \in \mathcal{A}$ for some $a' \in \text{ind}(\mathcal{A})$. But then, $B(a')$ has a derivation tree and thus $a' \in B^{\mathcal{I}}$ and $a \in (\exists r.B)^{\mathcal{I}}$.
- If $A \sqsubseteq \forall r.B \in \mathcal{T}$ and $a \in (\exists r^-.A)^{\mathcal{I}}$, there is some b with $b \in A^{\mathcal{I}}$ and $(b, a) \in r^{\mathcal{I}}$. In case $b \in \mathcal{A}$, $B(a)$ has a derivation tree by definition. If $b \notin \text{ind}(\mathcal{A})$, we get $a \in B^{\mathcal{I}_a}$. By (*), we know that $\mathcal{T}, \mathcal{A}_a \models B(a)$ and thus there $B(a)$ has a derivation tree. In both cases, the construction yields $a \in B^{\mathcal{I}}$.
- It remains to note that \mathcal{I} satisfies the functionality assertions from \mathcal{T} because of assumption (b). □

References

- Arenas, M., Botoeva, E., Calvanese, D., & Ryzhikov, V. (2016). Knowledge base exchange: The case of OWL 2 QL. *Artificial Intelligence*, 238, 11–62.
- Artale, A., Calvanese, D., Kontchakov, R., & Zakharyashev, M. (2009). The DL-Lite family and relations. *Journal of Artificial Intelligence Research*, 36, 1–69.
- Baader, F., Horrocks, I., Lutz, C., & Sattler, U. (2017). *An Introduction to Description Logics*. Cambridge University Press.
- Baget, J., Mugnier, M., Rudolph, S., & Thomazo, M. (2011). Walking the complexity lines for generalized guarded existential rules. In *Proceedings of International Conference on Artificial Intelligence (IJCAI)*, pp. 712–717.
- Bienvenu, M., Hansen, P., Lutz, C., & Wolter, F. (2016). First order-rewritability and containment of conjunctive queries in Horn description logics. In *Proceedings of International Conference on Artificial Intelligence (IJCAI)*, pp. 965–971.
- Bienvenu, M., Lutz, C., & Wolter, F. (2013). First order-rewritability of atomic queries in Horn description logics. In *Proceedings of International Conference on Artificial Intelligence (IJCAI)*, pp. 754–760.
- Bienvenu, M., & Ortiz, M. (2015). Ontology-mediated query answering with data-tractable description logics. In *Reasoning Web*, Vol. 9203 of *Lecture Notes in Computer Science*, pp. 218–307. Springer.
- Bienvenu, M., & Rosati, R. (2015). Query-based comparison of OBDA specifications. In *Proceedings of Description Logic Workshop (DL)*, Vol. 1350.

- Bienvenu, M., ten Cate, B., Lutz, C., & Wolter, F. (2014). Ontology-based data access: A study through disjunctive Datalog, CSP, and MMSNP. *ACM Transactions on Database Systems*, 39(4), 33:1–33:44.
- Botoeva, E., Konev, B., Lutz, C., Ryzhikov, V., Wolter, F., & Zakharyashev, M. (2016a). Inseparability and conservative extensions of description logic ontologies: A survey. In *Reasoning Web*, Vol. 9885 of *Lecture Notes in Computer Science*, pp. 27–89. Springer.
- Botoeva, E., Kontchakov, R., Ryzhikov, V., Wolter, F., & Zakharyashev, M. (2016b). Games for query inseparability of description logic knowledge bases. *Artificial Intelligence*, 234, 78–119.
- Botoeva, E., Lutz, C., Ryzhikov, V., Wolter, F., & Zakharyashev, M. (2019). Query inseparability for \mathcal{ALC} ontologies. *Artif. Intell.*, 272, 1–51.
- Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M., & Rosati, R. (2007). Tractable reasoning and efficient query answering in description logics: The DL-Lite family. *Journal of Automated Reasoning*, 39(3), 385–429.
- Comon, H., Dauchet, M., Gilleron, R., Löding, C., Jacquemard, F., Lugiez, D., Tison, S., & Tommasi, M. (2007). Tree automata techniques and applications. Available at <http://www.grappa.univ-lille3.fr/tata>. Release October, 12th 2007.
- Eiter, T., Gottlob, G., Ortiz, M., & Šimkus, M. (2008). Query answering in the description logic Horn- \mathcal{SHIQ} . In *Proceedings of European Conference on Logics in Artificial Intelligence (JELIA)*, Vol. 5293 of *LNCS*, pp. 166–179.
- Eiter, T., Ortiz, M., Šimkus, M., Tran, T.-K., & Xiao, G. (2012). Query rewriting for Horn- \mathcal{SHIQ} plus rules. In *Proceedings of AAAI Conference on Artificial Intelligence*.
- Emerson, E. A., & Jutla, C. S. (1991). Tree automata, mu-calculus and determinacy (extended abstract). In *Proceedings of the IEEE Annual Symposium Foundations of Computer Science (FOCS)*, pp. 368–377.
- Ghilardi, S., Lutz, C., & Wolter, F. (2006). Did I damage my ontology? A case for conservative extensions in description logics. In *Proceedings of International Conference on Principles Knowledge Representation and Reasoning (KR)*, pp. 187–197.
- Glimm, B., Lutz, C., Horrocks, I., & Sattler, U. (2008). Conjunctive query answering for the description logic \mathcal{SHIQ} . *Journal of Artificial Intelligence Research*, 31, 157–204.
- Grädel, E., & Walukiewicz, I. (1999). Guarded fixed point logic. In *Proceedings of Symposium on Logic in Computer Science (LICS)*, pp. 45–54.
- Grau, B. C., Horrocks, I., Kazakov, Y., & Sattler, U. (2008). Modular reuse of ontologies: Theory and practice. *Journal of Artificial Intelligence Research*, 31, 273–318.
- Gutiérrez-Basulto, V., Jung, J. C., & Sabellek, L. (2018). Reverse engineering queries in ontology-enriched systems: The case of expressive Horn description logic ontologies. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-ECAI-18)*.
- Horrocks, I., & Patel-Schneider, P. F. (2004). Reducing OWL entailment to description logic satisfiability. *Journal of Web Semantics*, 1(4), 345–357.

- Hustadt, U., Motik, B., & Sattler, U. (2007). Reasoning in description logics by a reduction to disjunctive Datalog. *Journal of Automated Reasoning*, 39(3).
- Ibáñez-García, Y., Lutz, C., & Schneider, T. (2014). Finite model reasoning in Horn description logics. In *Proceedings of International Conference on the Principles of Knowledge Representation and Reasoning (KR)*.
- Jung, J. C., Lutz, C., Martel, M., Schneider, T., & Wolter, F. (2017). Conservative extensions in guarded and two-variable fragments. In *Proceedings of International Colloquium on Automata, Languages, and Programming (ICALP)*, Vol. 80 of *LIPICs*, pp. 108:1–108:14.
- Jung, J. C., Papacchini, F., Wolter, F., & Zakharyashev, M. (2019). Model comparison games for Horn description logics. In *Proceedings of Symposium on Logic in Computer Science (LICS)*, pp. 1–14.
- Kazakov, Y. (2009). Consequence-driven reasoning for Horn-*SHIQ* ontologies. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 2040–2045.
- Konev, B., Ludwig, M., Walther, D., & Wolter, F. (2012). The logical difference for the lightweight description logic \mathcal{EL} . *Journal of Artificial Intelligence Research*, 44, 633–708.
- Konev, B., Lutz, C., Walther, D., & Wolter, F. (2009). Formal properties of modularisation. In Stuckenschmidt, H., Spaccapietra, S., & Parent, C. (Eds.), *Modular Ontologies*, Vol. 5445 of *LNCS*, pp. 25–66. Springer.
- Kontchakov, R., Pulina, L., Sattler, U., Schneider, T., Selmer, P., Wolter, F., & Zakharyashev, M. (2009). Minimal module extraction from DL-Lite ontologies using QBF solvers. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 836–840.
- Kontchakov, R., Wolter, F., & Zakharyashev, M. (2010). Logic-based ontology comparison and module extraction, with an application to DL-Lite. *Artif. Intell.*, 174, 1093–1141.
- Krötzsch, M., Rudolph, S., & Hitzler, P. (2013). Complexities of horn description logics. *ACM Trans. Comput. Log.*, 14(1), 2:1–2:36.
- Lutz, C. (2008). The complexity of conjunctive query answering in expressive description logics. In *Proceedings of International Conference on Automated Reasoning (IJCAR)*, Vol. 5195 of *LNCS*, pp. 179–193.
- Lutz, C., Seylan, I., & Wolter, F. (2012). An automata-theoretic approach to uniform interpolation and approximation in the description logic \mathcal{EL} . In *Proceedings of the International Conference on the Principles of Knowledge Representation and Reasoning (KR)*.
- Lutz, C., Walther, D., & Wolter, F. (2007). Conservative extensions in expressive description logics. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 453–458.
- Lutz, C., & Wolter, F. (2010). Deciding inseparability and conservative extensions in the description logic \mathcal{EL} . *Journal of Symbolic Computation*, 45(2), 194–228.
- Lutz, C., & Wolter, F. (2011). Foundations for uniform interpolation and forgetting in expressive description logics. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 989–995.

- Lutz, C., & Wolter, F. (2017). The data complexity of description logic ontologies. *Logical Methods in Computer Science*, 13(4).
- Matentzoglou, N., & Parsia, B. (2017). BioPortal Snapshot 30 March 2017 (data set). Zenodo research data repository. <http://doi.org/10.5281/zenodo.439510>.
- Ngo, N., Ortiz, M., & Šimkus, M. (2016). Closed predicates in description logics: Results on combined complexity. In *Proceedings of International Conference on the Principles of Knowledge Representation and Reasoning (KR)*, pp. 237–246.
- Poggi, A., Lembo, D., Calvanese, D., De Giacomo, G., Lenzerini, M., & Rosati, R. (2008). Linking data to ontologies. *Journal of Data Semantics*, 10, 133–173.
- Pratt, V. R. (1979). Models of program logics. In *Symposium on Foundations of Computer Science (FOCS)*, pp. 115–122.
- Rodriguez-Muro, M., & Calvanese, D. (2012). High performance query answering over DL-Lite ontologies. In *Proceedings of International Conference on Principles Knowledge Representation and Reasoning (KR)*.
- Tobies, S. (2001). *Complexity Results and Practical Algorithms for Logics in Knowledge Representation*. Ph.D. thesis, RWTH Aachen.
- Trivela, D., Stoilos, G., Chortaras, A., & Stamou, G. B. (2015). Optimising resolution-based rewriting algorithms for OWL ontologies. *Journal of Web Semantics*, 33, 30–49.
- Vardi, M. Y. (1998). Reasoning about the past with two-way automata. In *Proceedings of the International Colloquium on Automata, Languages, and Programming*, pp. 628–641.
- Wang, K., Wang, Z., Topor, R. W., Pan, J. Z., & Antoniou, G. (2014). Eliminating concepts and roles from ontologies in expressive descriptive logics. *Computational Intelligence*, 30(2), 205–232.
- Zhou, Y., Cuenca Grau, B., Nenov, Y., Kaminski, M., & Horrocks, I. (2015). PAGOdA: Pay-as-you-go ontology query answering using a Datalog reasoner. *Journal of Artificial Intelligence Research*, 54, 309–367.