

How to Approximate Ontology-Mediated Queries

Anneke Haga¹, Carsten Lutz¹, Leif Sabellek¹, Frank Wolter²

¹Department of Computer Science, University of Bremen, Germany

²Department of Computer Science, University of Liverpool, UK

{anneke,clu,sabellek}@uni-bremen.de, wolter@liverpool.ac.uk

Abstract

We introduce and study several notions of approximation for ontology-mediated queries based on the description logics \mathcal{ALC} and \mathcal{ALCI} . Our approximations are of two kinds: we may (1) replace the ontology with one formulated in a tractable ontology language such as \mathcal{ELI} or certain TGDs and (2) replace the database with one from a tractable class such as the class of databases whose treewidth is bounded by a constant. We determine the computational complexity and the relative completeness of the resulting approximations. (Almost) all of them reduce the data complexity from CONP -complete to PTIME , in some cases even to fixed-parameter tractable and to linear time. While approximations of kind (1) also reduce the combined complexity, this tends to not be the case for approximations of kind (2). In some cases, the combined complexity even increases.

1 Introduction

Ontology-mediated querying enriches database queries by an ontology, in this way providing domain knowledge and extending the language available for formulating queries. For ontologies written in popular expressive description logics (DLs) such as \mathcal{ALC} and \mathcal{ALCI} , however, the complexity of ontology-mediated querying is prohibitively high, CONP -complete in data complexity (Schaerf 1993) and EXPTIME - resp. 2EXPTIME -complete in combined complexity (Lutz 2008a). As a consequence, practical implementations often resort to approximating the answers to ontology mediated queries (OMQs) (Tserendorj et al. 2008; Thomas, Pan, and Ren 2010; Zhou et al. 2015), mostly using rather pragmatic approaches. The aim of this paper is to carry out a systematic study of OMQ approximation from a theoretical angle, introducing several principled notions of approximation and clarifying their computational complexity and relative completeness. In particular, we aim to find approximations that reduce the data complexity to PTIME or even to fixed-parameter tractability (FPT). Preferably, they should additionally reduce the combined complexity.

We mainly consider approximation from below, that is, approximations that are sound, but incomplete. While we also present some first results on approximation from above, whenever we speak of approximation without further qualification we mean approximation from below. An OMQ is a triple $Q = (\mathcal{O}, \Sigma, q)$ where \mathcal{O} is an ontology, q an actual

query such as a conjunctive query (CQ), and Σ a signature for the databases \mathcal{D} that Q is evaluated on. Our starting point is the observation that this gives us three points of attack for approximation: we can relax the ontology \mathcal{O} , the query q , and the database \mathcal{D} . However, relaxing the query is not useful for attaining PTIME data complexity as in the DLs mentioned above, ontology-mediated querying is CONP -hard already for atomic queries (AQs), that is, for very simple CQs of the form $A(x)$. We are thus left with the ontology and the database.

For ontology relaxing approximation, we choose an ontology language \mathcal{L} that admits ontology-mediated querying in PTIME in data complexity. To define approximate answers to OMQ $Q = (\mathcal{O}, \Sigma, q)$, we then consider all \mathcal{L} -ontologies \mathcal{O}' with $\mathcal{O} \models \mathcal{O}'$ (to guarantee soundness), replace \mathcal{O} with \mathcal{O}' , and take the union of the answers to the resulting OMQs Q' . Equivalently, we can use a single \mathcal{O}' , namely the unique *logically strongest* \mathcal{L} -ontology with $\mathcal{O} \models \mathcal{O}'$. Such \mathcal{O}' will typically be infinite (Bötcher, Lutz, and Wolter 2019; Haga et al. 2020), but it turns out that it never has to be materialized by an algorithm that computes approximate answers; the ontologies \mathcal{O}' in fact only serve the purpose of defining the semantics of approximation. As choices for \mathcal{L} , we consider Horn description logics such as \mathcal{ELI} and sets of restricted tuple-generating dependencies (TGDs), also known as existential rules and Datalog^\pm (Gottlob, Morak, and Pieris 2015). A related (but stronger) notion of OMQ approximation was proposed in (Haga et al. 2020).

For database relaxing approximation, we choose a class of databases \mathcal{D} that admits ontology-mediated querying in PTIME in data complexity. To define approximate answers to OMQ Q on input database \mathcal{D} , we then consider all databases $\mathcal{D}' \in \mathcal{D}$ such that there is a homomorphism from \mathcal{D}' to \mathcal{D} (to guarantee soundness) and take the union of the answers to Q on all such \mathcal{D}' . As choices for \mathcal{D} , we consider databases of bounded treewidth and databases that are proper trees. Equivalently, we can use the unraveling of \mathcal{D} into a structure of bounded treewidth, resp. into a tree. Such an unraveling may be infinite, but again this is unproblematic as the unraveling never has to be materialized by an algorithm that computes approximate answers.

Both of these approaches to approximation can also be used to define approximation from above. For ontology strengthening approximation, one requires $\mathcal{O}' \models \mathcal{O}$ and

for database strengthening approximation, one requires that there is a homomorphism from \mathcal{D} to \mathcal{D}' , rather than the other way around. In both cases, one then takes the intersection of the answers rather than the union. The resulting approximations are complete, but unsound. Note that approximation from above is particularly useful in combination with approximation from below (Tserendorj et al. 2008; Zhou et al. 2015). If both approximations produce the same answers, one has actually succeeded to compute the ‘real’, non-approximate answers.

An OMQ language is a pair $(\mathcal{L}, \mathcal{Q})$ with \mathcal{L} an ontology language and \mathcal{Q} a query language. We consider the approximation of OMQ languages $(\mathcal{L}, \mathcal{Q})$ where $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCI}\}$ and $\mathcal{Q} \in \{\text{UCQ}, \text{CQ}, \text{AQ}, \text{bELIQ}\}$ with UCQ denoting unions of CQs and bELIQ denoting the class of unary CQs that correspond to \mathcal{ELI} -concepts (ELIQs) and of Boolean CQs $\exists x q(x)$ with $q(x)$ an ELIQ. The exact problem that we consider is (approximate) *OMQ evaluation*, meaning to decide, given an OMQ Q , a database \mathcal{D} , and a tuple \bar{a} of constants from \mathcal{D} , whether \bar{a} is an (approximate) answer to Q on \mathcal{D} . We give an overview of our results.

We start in Section 4 with ontology relaxing approximation, choosing for \mathcal{L} the description logic \mathcal{ELI}_{\perp}^u , that is, the extension of \mathcal{ELI} with the universal role and bottom (logical falsity). We then prove that \mathcal{ELI}_{\perp}^u -ontology relaxing OMQ evaluation is in PTIME in data complexity and EXPTIME-complete in combined complexity for all OMQ languages $(\mathcal{L}, \mathcal{Q})$ mentioned above. In contrast, non-approximate OMQ evaluation is 2EXPTIME-complete in $(\mathcal{ALCI}, \mathcal{Q})$ for $\mathcal{Q} \in \{\text{CQ}, \text{UCQ}\}$. If we consider more restricted classes of queries, the complexity improves further. In fact, \mathcal{ELI}_{\perp}^u -ontology relaxing OMQ evaluation is fixed-parameter tractable (FPT) with single exponential overall running time if \mathcal{Q} is the class of all CQs or UCQs whose treewidth is bounded by a constant; here and in what follows, the parameter is the size of the OMQ. For $\mathcal{Q} = \text{bELIQ}$, we even obtain linear time in data complexity.

In Section 5, we consider tree-database relaxing approximation. These (almost) deliver the same answers as \mathcal{ELI}_{\perp}^u -ontology relaxing approximation in $(\mathcal{ALCI}, \text{bELIQs})$, but are incomparable for more expressive query languages. They turn out to be less well-behaved regarding combined complexity, being 2EXPTIME-complete in $(\mathcal{ALCI}, \mathcal{Q})$ for $\mathcal{Q} \in \{\text{UCQ}, \text{CQ}\}$; this in fact even holds for (U)CQs of bounded treewidth. If $\mathcal{L} = \mathcal{ALC}$ or $\mathcal{Q} \in \{\text{AQ}, \text{bELIQ}\}$, then they are only EXPTIME-complete, as in the non-approximate case. On the other hand, tree-database relaxing approximation enjoys a slight advantage in data complexity over \mathcal{ELI}_{\perp}^u -ontology relaxing approximation, namely linear time for *all* OMQ languages $(\mathcal{Q}, \mathcal{L})$ considered in this paper. We also prove the surprising result that tree-database relaxing OMQ evaluation is EXPSpace-hard in $(\mathcal{ALC}, \text{CQ})$ and 2EXPTIME-complete in $(\mathcal{ALC}, \text{UCQ})$. This means that it is *harder* in combined complexity than non-approximate OMQ evaluation (which is only EXPTIME-complete in these two cases), while it is easier in data complexity.

In Section 6, we revisit ontology relaxing approximation, replacing \mathcal{ELI}_{\perp}^u with frontier-one TGDs whose rule bodies and heads are of bounded treewidth. Recall that a TGD is

frontier-one if body and head share at most a single variable (Baget et al. 2009). For rule heads, we do not only require bounded treewidth, but the existence of a tree decomposition in which the bags overlap in at most a *single* element. The resulting approximations are significantly more complete than \mathcal{ELI}_{\perp}^u -ontology relaxing approximations, but enjoy the same favourable computational properties regarding both data and combined complexity except that we do not attain linear time. We also observe that by increasing the treewidth of the rule bodies and heads, we obtain infinite hierarchies of increasingly complete approximations.

In Section 7, we generalize tree-database relaxing approximation into btw-database relaxing approximation, replacing tree databases with databases of bounded treewidth. These are strictly more complete than TGD-ontology relaxing approximations. They enjoy the same computational properties as tree-database relaxing approximations both regarding data and combined complexity except that we do not attain linear time.

Finally, in Section 8 we consider approximation from above. For database strengthening approximation, the results are negative: we show CONP-completeness in data complexity even if the original OMQ is from $(\mathcal{EL}, \text{CQ})$, an OMQ language that admits non-approximate OMQ evaluation in PTIME in data complexity. Ontology strengthening approximation looks more promising. We consider the fragment \mathcal{ELIU}_{\perp} of \mathcal{ALC} and show that \mathcal{ELIU}_{\perp} -ontology relaxing OMQ evaluation in $(\mathcal{ELIU}_{\perp}, \mathcal{Q})$ is FPT with double exponential overall running time (thus in PTIME in data complexity) and 2EXPTIME-complete in combined complexity for $\mathcal{Q} \in \{\text{AQ}, \text{CQ}, \text{UCQ}\}$. Note that non-approximate OMQ evaluation in $(\mathcal{ELIU}_{\perp}, \text{AQ})$ is only EXPTIME-complete (Lutz 2008a), and thus this is another case where approximate OMQ evaluation is harder in combined complexity than non-approximate OMQ evaluation. An appendix with full proofs is available at (Haga et al. 2021).

Related work. Several approaches achieve practically efficient OMQ evaluation by a pragmatic translation of ontologies into languages that enjoy PTIME data complexity such as Datalog, OWL 2 QL, or OWL 2 EL. This includes Screech (Hitzler and Vrandečić 2005; Tserendorj et al. 2008), TrOWL (Pan and Thomas 2007; Thomas, Pan, and Ren 2010), and PAGODA (Zhou et al. 2015), see also (Carral et al. 2014). Approximations are computed both from below and above, in the spirit of knowledge compilation (Selman and Kautz 1996). Approximations of ontologies in tractable languages with stronger guarantees were recently studied in (Bötcher, Lutz, and Wolter 2019; Haga et al. 2020), but querying and data were (mostly) not considered. In database theory, approximate querying (without ontologies) was studied in (Fink and Olteanu 2011; Barceló, Libkin, and Romero 2014; Barceló, Romero, and Zeume 2018). The approximation and rewriting of OMQs in(to) datalog is studied in (Bienvenu et al. 2014; Kaminski, Nenov, and Grau 2016; Feier, Kuusisto, and Lutz 2019). In the context of Horn DLs, OMQ approximation that achieves FPT was considered in (Barceló et al. 2019).

2 Preliminaries

Description Logics. Let \mathbb{N}_C and \mathbb{N}_R be countably infinite sets of *concept names* and *role names*. Further fix a countably infinite supply of *constants*. A *role* is a role name r or an *inverse role* r^- , with r a role name and $(r^-)^- = r$. An *ALCI-concept* is defined according to the syntax rule

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

where A ranges over concept names and r over roles. We use $C \sqcup D$ as abbreviation for $\neg(\neg C \sqcap \neg D)$, $C \rightarrow D$ for $\neg C \sqcup D$, and $\forall r.C$ for $\neg \exists r.\neg C$. An \mathcal{ELI}_\perp -concept is an \mathcal{ALCI} -concept that does not use negation “ \neg ” and an \mathcal{ELI}_\perp^u -concept is an \mathcal{ELI}_\perp -concept that may use the *universal role* u in place of a role name. \mathcal{ELI} -concepts and \mathcal{ELI}^u -concepts do not admit \perp . Let $\mathcal{L} \in \{\mathcal{ALCI}, \mathcal{ELI}, \mathcal{ELI}^u, \mathcal{ELI}_\perp, \mathcal{ELI}_\perp^u\}$. An \mathcal{L} -ontology is a finite set of *concept inclusion (CIs)* $C \sqsubseteq D$ with C and D \mathcal{L} -concepts. We sometimes also consider infinite ontologies that, however, only serve the purpose of defining a semantics and never have to be represented explicitly. For \mathcal{ELI}_\perp^u -ontologies, we assume w.l.o.g. that \perp occurs only in CIs of the form $C \sqsubseteq \perp$ where \perp does not occur in C . A *database* is a finite set of *facts* of the form $A(a)$ or $r(a, b)$ where $A \in \mathbb{N}_C \cup \{\top\}$, $r \in \mathbb{N}_R$, and a, b are constants. We use $\text{adom}(\mathcal{D})$ to denote the set of constants used in database \mathcal{D} , also called its *active domain*.

A *signature* Σ is a set of concept and role names, uniformly referred to as *symbols*. We use $\text{sig}(X)$ to denote the set of symbols used in any syntactic object X such as a concept or an ontology. A Σ -database is a database \mathcal{D} with $\text{sig}(\mathcal{D}) \subseteq \Sigma$. The *size* of a (finite) syntactic object X , denoted $\|X\|$, is the number of symbols needed to write it as a word using a suitable encoding.

The semantics is given in terms of *interpretations* \mathcal{I} , which we define to be a (possibly infinite and) non-empty set of facts. We use $\Delta^\mathcal{I}$ to denote the set of individual names in \mathcal{I} , define $A^\mathcal{I} = \{a \mid A(a) \in \mathcal{I}\}$ for all $A \in \mathbb{N}_C$, and $r^\mathcal{I} = \{(a, b) \mid r(a, b) \in \mathcal{I}\}$ for all $r \in \mathbb{N}_R$. The extension $C^\mathcal{I}$ of \mathcal{ALCI} -concepts C is then defined as usual (Baader et al. 2017). The universal role u is always interpreted as $u^\mathcal{I} = \Delta^\mathcal{I} \times \Delta^\mathcal{I}$. This definition of interpretation is slightly different from the usual one, but equivalent; its virtue is uniformity as every database is a (finite) interpretation. Interpretation \mathcal{I} *satisfies* CI $C \sqsubseteq D$ if $C^\mathcal{I} \subseteq D^\mathcal{I}$, fact $A(a)$ if $a \in A^\mathcal{I}$, and fact $r(a, b)$ if $(a, b) \in r^\mathcal{I}$. We thus make the standard names assumption, that is, we interpret constants as themselves. For $S \subseteq \Delta^\mathcal{I}$, we use $\mathcal{I}|_S$ to denote the restriction of \mathcal{I} to facts that only contain constants from S .

Interpretation \mathcal{I} is a *model* of an ontology or database if it satisfies all inclusions or facts in it. A database \mathcal{D} is *satisfiable* w.r.t. an ontology \mathcal{O} if there is a model \mathcal{I} of \mathcal{O} and \mathcal{D} . We write $\mathcal{O} \models \mathcal{O}'$ if every model of ontology \mathcal{O} is also a model of ontology \mathcal{O}' . We associate every interpretation \mathcal{I} with an undirected graph $G_\mathcal{I} = (V, E)$ where $V = \Delta^\mathcal{I}$ and $E = \{\{d, e\} \mid (d, e) \in r^\mathcal{I} \text{ for some } r \in \mathbb{N}_R\}$. We say that \mathcal{I} is a *tree* if $G_\mathcal{I}$ is acyclic without self loops and multi-edges, that is, $(d, e) \in r_1^\mathcal{I}$ implies $(d, e) \notin r_2^\mathcal{I}$ for all distinct roles r_1, r_2 . Note that, somewhat unusually, our trees need thus not be connected.

Queries. A *conjunctive query (CQ)* is of the form $q(\bar{x}) = \exists \bar{y} \varphi(\bar{x}, \bar{y})$, where \bar{x} and \bar{y} are tuples of variables and $\varphi(\bar{x}, \bar{y})$ is a conjunction of *atoms* of the form $A(x)$ and $r(x, y)$, $A \in \mathbb{N}_C$, $r \in \mathbb{N}_R$, and x, y variables from $\bar{x} \cup \bar{y}$. We require that all variables in \bar{x} are used in φ , call the variables in \bar{x} the *answer variables* of q , and use $\text{var}(q)$ to denote $\bar{x} \cup \bar{y}$. We take the liberty to write $\alpha \in q$ to indicate that α is an atom in q and sometimes write $r^-(x, y) \in q$ in place of $r(y, x) \in q$. The CQ q gives rise to a database \mathcal{D}_q , often called the canonical database for q , obtained by viewing the variables in q as constants and the atoms as facts. For $V \subseteq \text{var}(q)$, we use $q|_V$ to denote the restriction of q to the atoms that use only variables in V .

A *homomorphism* from interpretation \mathcal{I}_1 to interpretation \mathcal{I}_2 is a function $h : \Delta^{\mathcal{I}_1} \rightarrow \Delta^{\mathcal{I}_2}$ such that $d \in A^{\mathcal{I}_1}$ implies $h(d) \in A^{\mathcal{I}_2}$ and $(d, e) \in r^{\mathcal{I}_1}$ implies $(h(d), h(e)) \in r^{\mathcal{I}_2}$ for all $d, e \in \Delta^{\mathcal{I}_1}$, $A \in \mathbb{N}_C$, and $r \in \mathbb{N}_R$. A homomorphism from CQ q to interpretation \mathcal{I} is a homomorphism from \mathcal{D}_q to \mathcal{I} . A tuple $\bar{d} \in (\Delta^\mathcal{I})^{|\bar{x}|}$ is an *answer* to q on \mathcal{I} if there is a homomorphism h from q to \mathcal{I} with $h(\bar{x}) = \bar{d}$. A *contraction* of a CQ p is a CQ that can be obtained from q by identifying variables. The identification of two answer variables is not admitted and the identification of an answer variable x with a quantified variable y results in x .

A *union of conjunctive queries (UCQ)* $q(\bar{x})$ is a disjunction of CQs that all have the same answer variables \bar{x} . A tuple $\bar{d} \in (\Delta^\mathcal{I})^{|\bar{x}|}$ is an *answer* to q on interpretation \mathcal{I} , written $\mathcal{I} \models q(\bar{d})$, if \bar{d} is an answer to some CQ in q on \mathcal{I} . We use $q(\mathcal{I})$ to denote set of all answers to q on \mathcal{I} . The *arity* of q is the length of \bar{x} and q is *Boolean* if it is of arity zero.

An \mathcal{ELI} -query (*ELIQ*) is a unary CQ $q(x)$ such that \mathcal{D}_q is a connected tree and a *Boolean \mathcal{ELI} -query (BELIQ)* is a Boolean CQ $q()$ such that \mathcal{D}_q is a connected tree. One can alternatively define ELIQs as being of the form $C(x)$ with C an \mathcal{ELI} -concept, and BELIQs as being of the form $\exists u.C$ with C an \mathcal{ELI} -concept and we may thus use ELIQs as \mathcal{ELI} -concepts and BELIQs as \mathcal{ELI}^u -concepts, and vice versa. For uniformity, we use *BELIQ* to refer to a CQ that is either an ELIQ or a BELIQ. An *atomic query (AQ)* is an ELIQ of the form $A(x)$, A a concept name.

Ontology-Mediated Queries. An *ontology-mediated query (OMQ)* is a triple $Q = (\mathcal{O}, \Sigma, q)$ with \mathcal{O} an ontology, $\Sigma \subseteq \text{sig}(\mathcal{O}) \cup \text{sig}(q)$ a signature called the *data signature*, and q a query such as a UCQ. We write $Q(\bar{x})$ to indicate that the answer variables of q are \bar{x} . The signature Σ expresses the promise that Q is only evaluated on Σ -databases. Let \mathcal{D} be such a database. Then $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$ is an *answer* to Q on \mathcal{D} , written $\mathcal{D} \models Q(\bar{a})$, if $\mathcal{I} \models q(\bar{a})$ for all models \mathcal{I} of \mathcal{O} and \mathcal{D} . When more convenient, we might alternatively write $\mathcal{D}, \mathcal{O} \models q(\bar{a})$. We further write $Q(\mathcal{D})$ to denote the set of all answers to Q on \mathcal{D} . For OMQs $Q_1(\bar{x})$ and $Q_2(\bar{x})$, $Q_i = (\mathcal{O}_i, \Sigma, q_i)$, we say that Q_1 is *contained* in Q_2 and write $Q_1 \subseteq Q_2$, if for every Σ -database \mathcal{D} , $Q_1(\mathcal{D}) \subseteq Q_2(\mathcal{D})$. We say that Q_1 is *equivalent* to Q_2 and write $Q_1 \equiv Q_2$, if $Q_1 \subseteq Q_2$ and $Q_2 \subseteq Q_1$. We use $(\mathcal{L}, \mathcal{Q})$ to denote the *OMQ language* that contains all OMQs Q in which \mathcal{O} is formulated in DL \mathcal{L} and q in query language \mathcal{Q} , such as in $(\mathcal{ALCI}, \text{UCQ})$ and $(\mathcal{ELI}, \text{AQ})$.

Treewidth. Treewidth is a widely used notion that measures the degree of tree-likeness of a graph. As common for example in the area of constraint satisfaction problems, we are interested in two parameters of tree decompositions instead of only one. A *tree decomposition* of an interpretation \mathcal{I} is a triple $(V, E, (B_v)_{v \in V})$ where (V, E) is an undirected tree and $(B_v)_{v \in V}$ is a family of subsets of $\Delta^{\mathcal{I}}$, often referred to as bags, such that:

1. for all $d \in \Delta^{\mathcal{I}}$, $\{v \in V \mid d \in B_v\}$ is nonempty and connected in (V, E) ;
2. if $(d_1, d_2) \in r^{\mathcal{I}}$ for any role name r , then there is a $v \in V$ with $d_1, d_2 \in B_v$.

We call $(V, E, (B_v)_{v \in V})$ an (ℓ, k) -*tree decomposition* if for all distinct $v, v' \in V$, $|B_v \cap B_{v'}| \leq \ell$ and $|B_v| \leq k$. An interpretation \mathcal{I} has *treewidth* (ℓ, k) if it admits an (ℓ, k) -tree decomposition. It has *treewidth* k if it has treewidth $(k, k + 1)$. As usual, the ‘+1’ is used to achieve that trees have treewidth 1.

We also speak of the treewidth of a CQ $q = \exists \bar{y} \varphi(\bar{x}, \bar{y})$, which is that of $\mathcal{D}_q|_{\bar{y}}$, and of the treewidth of a UCQ q , which is the maximum of the treewidths of the CQs in q . Note that answer variables do not contribute to treewidth. For $\ell, k \geq 1$ with $\ell < k$, we use $\text{CQ}_{\ell, k}^{\text{tw}}$ (resp. CQ_k^{tw}) to denote the class of CQs of treewidth (ℓ, k) (resp. of treewidth k), and likewise for $\text{UCQ}_{\ell, k}^{\text{tw}}$ (resp. UCQ_k^{tw}) and UCQs.

Tuple-Generating Dependencies. A *tuple-generating dependency* (TGD) is a first-order sentence ϑ of the form $\forall \bar{x} \forall \bar{y} (\phi(\bar{x}, \bar{y}) \rightarrow \exists \bar{z} \psi(\bar{x}, \bar{z}))$ such that $\exists \bar{y} \phi(\bar{x}, \bar{y})$ and $\exists \bar{z} \psi(\bar{x}, \bar{z})$ are CQs. For simplicity, we write ϑ as $\phi(\bar{x}, \bar{y}) \rightarrow \exists \bar{z} \psi(\bar{x}, \bar{z})$. We call ϕ and ψ the *body* and *head* of ϑ . The body may be the empty conjunction, i.e. logical truth, then denoted by \top , and the head may be logical falsity denoted by \perp . TGDs with head \perp are often called *denial constraints* (Fan and Geerts 2012). The *frontier variables* of TGD ϑ are the variables that occur unquantified in both body and head. We say that ϑ is *frontier one* if it has at most one frontier variable. An interpretation \mathcal{I} *satisfies* ϑ , denoted $\mathcal{I} \models \vartheta$, if $q_\phi(\mathcal{I}) \subseteq q_\psi(\mathcal{I})$.

We also consider ontologies that are sets of TGDs and, more generally, sentences formulated in first-order logic (FO). What we mean here is the version of FO in which only unary and binary relation symbols are used, which are from N_C and N_R , respectively. Function symbols, constants, and equality are not permitted. An interpretation is a *model* of an FO ontology if it satisfies all sentences in \mathcal{O} . With an *FO-fragment* \mathcal{L} , we mean a class of FO-sentences and an *\mathcal{L} -ontology* is a finite set of sentences from \mathcal{L} . As in the DL case, we sometimes also consider infinite ontologies. We use TGD to denote the FO-fragment that consists of all TGDs, and thus speak of *TGD-ontologies*. It is easy to see that every \mathcal{ELI}_\perp^u -ontology is also a frontier one TGD-ontology.

A standard tool for dealing with TGD ontologies \mathcal{O} is the *chase* that constructs from \mathcal{O} and a database \mathcal{D} a universal model of \mathcal{O} and \mathcal{D} , that is, a model $\text{ch}_{\mathcal{O}}(\mathcal{D})$ of \mathcal{D} and \mathcal{O} that homomorphically embeds into every model of \mathcal{D} and \mathcal{O} and thus satisfies $\mathcal{D}, \mathcal{O} \models q(\bar{a})$ iff $\text{ch}_{\mathcal{O}}(\mathcal{D}) \models q(\bar{a})$ for all CQs q and tuples \bar{a} . Details are given in the appendix, see also (Johnson and Klug 1982; Cali, Gottlob, and Kifer 2013).

3 OMQ Approximation

We introduce two notions of OMQ approximation from below: one where we relax the ontology and one where we relax the database. We start with the former.

For an OMQ $Q(\bar{x}) = (\mathcal{O}, \Sigma, q)$, a Σ -database \mathcal{D} , and an ontology language \mathcal{L}' , we use $\text{app}_{\mathcal{L}'}(Q, \mathcal{D})$ to denote the set of tuples $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$ such that $\bar{a} \in Q'(\mathcal{D})$ for some OMQ $Q' = (\mathcal{O}', \Sigma, q)$ where \mathcal{O}' is a (finite) \mathcal{L}' -ontology with $\mathcal{O} \models \mathcal{O}'$. Note that the ontology \mathcal{O}' might contain symbols that do not occur in \mathcal{O} , we will see later that this in fact results in additional answers. Every choice of an OMQ language $(\mathcal{L}, \mathcal{Q})$ and an ontology language \mathcal{L}' gives rise to an approximate OMQ evaluation problem, as follows.

\mathcal{L}' -ontology relaxing OMQ evaluation in $(\mathcal{L}, \mathcal{Q})$
INPUT: OMQ $Q(\bar{x}) = (\mathcal{O}, \Sigma, q) \in (\mathcal{L}, \mathcal{Q})$,
 Σ -database \mathcal{D} , tuple $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$
OUTPUT: ‘yes’ if $\bar{a} \in \text{app}_{\mathcal{L}'}(Q, \mathcal{D})$ and ‘no’ otherwise

It follows from the definition that ontology relaxing approximation is sound, in the sense that $\text{app}_{\mathcal{L}'}(Q, \mathcal{D}) \subseteq Q(\mathcal{D})$ for all OMQs $Q = (\mathcal{O}, \Sigma, q)$ and Σ -databases \mathcal{D} . We concentrate on the case where \mathcal{L}' is an ontology language such that $(\mathcal{L}', \mathcal{Q})$ enjoys PTIME data complexity. For $(\mathcal{L}, \mathcal{Q}) = (\mathcal{ALCCl}, \text{CQ})$, for instance, we might choose $\mathcal{L}' = \mathcal{ELI}_\perp^u$.

In the definition of $\text{app}_{\mathcal{L}'}(Q, \mathcal{D})$, one can equivalently replace the infinitely many \mathcal{L}' -ontologies \mathcal{O}' with the single but infinite \mathcal{L}' -ontology $\mathcal{O}_{\mathcal{L}'}$, that consists of all \mathcal{L}' -sentences φ with $\mathcal{O} \models \varphi$. In fact, the following lemma is a consequence of compactness. For an OMQ $Q(\bar{x}) = (\mathcal{O}, \Sigma, q)$ and an ontology language \mathcal{L}' , we use $Q_{\mathcal{L}'}^{\approx}$ to denote the OMQ $(\mathcal{O}_{\mathcal{L}'}, \Sigma, q)$.

Lemma 1. *Let $Q(\bar{x}) = (\mathcal{O}, \Sigma, q) \in (\text{FO}, \text{UCQ})$ be an OMQ and \mathcal{L}' an ontology language. Then for every Σ -database \mathcal{D} , $\text{app}_{\mathcal{L}'}(Q, \mathcal{D}) = Q_{\mathcal{L}'}^{\approx}(\mathcal{D})$.*

Note that we do not insist that the infinitely many \mathcal{O}' or the infinite $\mathcal{O}_{\mathcal{L}'}$ is ever explicitly generated when computing approximate answers.

We next define a mode of approximation that is based on relaxing the database by replacing it with a homomorphic pre-image. A *pointed database* is a pair (\mathcal{D}, \bar{b}) with \mathcal{D} a database and \bar{b} a tuple over $\text{adom}(\mathcal{D})$. For an OMQ $Q(\bar{x}) = (\mathcal{O}, \Sigma, q)$, a Σ -database \mathcal{D} , and a class of pointed databases \mathfrak{D} , we use $\text{app}_{\mathfrak{D}}(Q, \mathcal{D})$ to denote the set of tuples $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$ such that for some $(\mathcal{D}', \bar{b}) \in \mathfrak{D}$ with $\bar{b} \in Q(\mathcal{D}')$ there is a homomorphism h from \mathcal{D}' to \mathcal{D} with $h(\bar{b}) = \bar{a}$. Every choice of $(\mathcal{L}, \mathcal{Q})$ and \mathfrak{D} gives rise to an approximate OMQ evaluation problem, as follows.

\mathfrak{D} -database relaxing OMQ evaluation in $(\mathcal{L}, \mathcal{Q})$
INPUT: OMQ $Q(\bar{x}) = (\mathcal{O}, \Sigma, q) \in (\mathcal{L}, \mathcal{Q})$,
 Σ -database \mathcal{D} , tuple $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$
OUTPUT: ‘yes’ if $\bar{a} \in \text{app}_{\mathfrak{D}}(Q, \mathcal{D})$ and ‘no’ otherwise

Answers to any OMQ $Q = (\mathcal{O}, \Sigma, q) \in (\text{FO}, \text{UCQ})$ are preserved under homomorphisms if \mathcal{O} does not use equality, that is, if $\mathcal{D}_1, \mathcal{D}_2$ are databases, h is a homomorphism from \mathcal{D}_1 to \mathcal{D}_2 , and $\bar{a} \in Q(\mathcal{D}_1)$ for an OMQ Q , then $h(\bar{a}) \in$

$Q(\mathcal{D}_2)$ (Bienvenu et al. 2014). As a consequence, database relaxing approximation is sound.

We are interested in choosing \mathfrak{D} such that evaluating OMQs from $(\mathcal{L}, \mathcal{Q})$ on \mathfrak{D} enjoys PTIME data complexity. An important example are classes of databases that are of bounded treewidth, the simplest case being the class of databases that are trees. More precisely, we use \mathfrak{D}_λ (with ‘ λ ’ symbolizing a tree) to denote the class of all pointed databases (\mathcal{D}, \bar{a}) such that the restriction of \mathcal{D} to domain $\text{adom}(\mathcal{D}) \setminus \bar{a}$ is a tree. Recall that a tree does not need to be connected. The resulting notion of approximation is closely related to \mathcal{ELT}_\perp^u -ontology relaxing approximation.

In the same way in which we have rephrased ontology relaxing approximation in terms of a single infinite ontology, we can sometimes (depending on the choice of \mathfrak{D}) rephrase database relaxing approximation in terms of evaluation on a single infinite database. We illustrate this for the case $\mathfrak{D} = \mathfrak{D}_\lambda$. Let \mathcal{D} be a database and $S \subseteq \text{adom}(\mathcal{D})$. A path in \mathcal{D} is a sequence $p = a_0 r_1 a_1 r_2 \dots r_n a_n$, $n \geq 0$, where $a_0, \dots, a_n \in \text{adom}(\mathcal{D})$, r_1, \dots, r_n are (potentially inverse) roles, and $r_{i+1}(a_i, a_{i+1}) \in \mathcal{D}$ for $0 \leq i < n$. We use $\text{tail}(p)$ to denote a_n . The tree unraveling \mathcal{D}_S^\approx of \mathcal{D} at S is the (potentially infinite) database that contains the following facts: all facts from $\mathcal{D}|_S$, $r(p, prb)$ for every path prb , $A(p)$ for every path p with $A(\text{tail}(p)) \in \mathcal{D}$, and $r(a, p)$ for every $r(a, b) \in \mathcal{D}$ and every path p with $a \in S$ and $\text{tail}(p) = b$ (r a potentially inverse role). Note that \mathcal{D}_S^\approx is a tree if and only if $S = \emptyset$. Thus, the tree unravelings $\mathcal{D}_{\bar{a}}^\approx$ in the following lemma are in general not tree databases.

Lemma 2. *Let $Q = (\mathcal{O}, \Sigma, q) \in (FO, UCQ)$. Then for all Σ -databases \mathcal{D} and $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{a}|}$, $\bar{a} \in \text{app}_{\mathfrak{D}_\lambda}(Q, \mathcal{D})$ iff $\bar{a} \in Q(\mathcal{D}_{\bar{a}}^\approx)$.*

4 \mathcal{ELT}_\perp^u -Ontology Relaxing Approximation

We consider \mathcal{ELT}_\perp^u -ontology relaxing evaluation of OMQs from (\mathcal{ALCC}, UCQ) , starting with an example.

Example 1. *Let $Q(x) = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALCC}, CQ)$ where*

$$\begin{aligned} \mathcal{O} &= \{\top \sqsubseteq \forall r.(B_1 \rightarrow A) \sqcup \forall r.(B_2 \rightarrow A)\} \\ \Sigma &= \{r, A, B_1, B_2\} \\ q &= \exists y r(x, y) \wedge A(y). \end{aligned}$$

Further let $\mathcal{D} = \{r(a, b_1), r(a, b_2), B_1(b_1), B_2(b_2)\}$. Clearly, $a \in Q(\mathcal{D})$. The ontology $\mathcal{O}_{\mathcal{ELT}_\perp^u}^\approx$ contains CI $\exists r.B_1 \sqcap \exists r.B_2 \sqsubseteq \exists r.A$, thus also $a \in Q_{\mathcal{ELT}_\perp^u}^\approx(\mathcal{D})$.

We next illustrate incompleteness, which cannot be avoided by any notion of approximation from below that attains PTIME data complexity. This follows from the existence of OMQs that express non-3-colorability.

Example 2. *Let $Q() = (\mathcal{O}, \Sigma, \exists x D(x)) \in (\mathcal{ALCC}, BELIQ)$ with*

$$\mathcal{O} = \{\top \sqsubseteq R \sqcup G \sqcup B, X \sqcap \exists e.X \sqsubseteq D \mid X \in \{R, G, B\}\}$$

and $\Sigma = \{e\}$. Every Σ -database \mathcal{D} can be viewed as an undirected graph by ‘forgetting’ the direction of e -edges. Then $\mathcal{D} \models Q$ iff the graph is not 3-colorable. In contrast, the careful chase algorithm given below can be used to verify

that $\mathcal{D} \not\models Q_{\mathcal{ELT}_\perp^u}^\approx$ for all Σ -databases \mathcal{D} . It is easy to modify the example so as to use an AQ in place of a BELIQ.

For readers who are disappointed by the extreme incompleteness in the previous example, we remark replacing \mathcal{ELT}_\perp^u with classes of TGDs improves the situation.

We next present three observations regarding our definition of ontology relaxing approximation. The first observation is that it increases completeness to admit in $\mathcal{O}_{\mathcal{ELT}_\perp^u}^\approx$ symbols that do not occur in \mathcal{O} .

Example 3. *Let $Q(x) = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALCC}, CQ)$ with*

$$\begin{aligned} \mathcal{O} &= \{\top \sqsubseteq \forall r.(B_1 \rightarrow B) \sqcup \forall r.(B_2 \rightarrow B)\} \\ \Sigma &= \{r, A, B, B_1, B_2\} \quad q = \exists y r(x, y) \wedge A(y) \wedge B(y). \end{aligned}$$

Then $\mathcal{O}_{\mathcal{ELT}_\perp^u}^\approx$ contains the CI

$$\exists r.(B_1 \sqcap A) \sqcap \exists r.(B_2 \sqcap A) \sqsubseteq \exists r.(A \sqcap B)$$

despite the fact that A does not occur in \mathcal{O} . Let

$$\mathcal{D} = \{r(a, b_1), r(a, b_2), B_1(b_1), B_2(b_2), A(b_1), A(b_2)\}.$$

Then $a \in Q_{\mathcal{ELT}_\perp^u}^\approx(\mathcal{D})$, but we show in the appendix that this is no longer true when we remove from $\mathcal{O}_{\mathcal{ELT}_\perp^u}^\approx$ all CIs that use a symbol that does not occur in \mathcal{O} .

It is, however, easy to see that it suffices to admit in $\mathcal{O}_{\mathcal{ELT}_\perp^u}^\approx$ the symbols that occur in \mathcal{O} or in Σ while additional symbols do not further increase completeness.

The second observation is that \mathcal{ELT}_\perp^u -ontology relaxing approximation is more complete than \mathcal{ELT}_\perp -ontology relaxing approximation. In fact, it seems to be much more challenging to compute answers for the latter while offering no obvious benefit compared to the former, and thus we do not consider it in this paper.

Example 4. *Let $Q() = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALCC}, CQ)$ be the Boolean OMQ with*

$$\mathcal{O} = \{A \sqsubseteq B \sqcup \forall r.B\} \quad \Sigma = \{r, A, B\} \quad q = \exists x B(x).$$

Let $\mathcal{D} = \{A(a), r(a, b)\}$. Then $\mathcal{O}_{\mathcal{ELT}_\perp^u}^\approx$ contains $A \sqcap \exists r.\top \sqsubseteq \exists u.B$ and thus $\mathcal{D} \models Q_{\mathcal{ELT}_\perp^u}^\approx$, but it is shown in the appendix that $\mathcal{D} \not\models Q_{\mathcal{ELT}_\perp}^\approx$.

We remark that \mathcal{ELT}_\perp^u -ontology relaxing approximation is also more complete than \mathcal{ELT}_\perp^u -ontology relaxing approximation, examples are easy to find.

Third, we observe that defining ontology relaxing approximation in terms of ontologies that are implied by the original ontology does not necessarily result in maximum completeness. In fact, the following example shows that it may pay off to use an \mathcal{ELT}_\perp^u -ontology that is *not* a consequence of the original ontology. This is a very interesting effect, but we do not investigate it further.

Example 5. *We use the ontology \mathcal{O} and signature Σ from Example 4. Let $Q_0() = (\mathcal{O}, \Sigma, q_0)$ for*

$$q_0 = \exists x \exists y A(x) \wedge A(y) \wedge B(y) \wedge r(x, y) \wedge r(y, x),$$

and let $\mathcal{D}_0 = \{A(a), r(a, b), r(b, a), A(b)\}$. While $\mathcal{D}_0 \models Q_0$, the careful chase algorithm below yields $\mathcal{D}_0 \not\models Q_{\mathcal{ELT}_\perp^u}^\approx$. Now take $\mathcal{O}' = \{A \sqsubseteq B\}$ and note that $\mathcal{O} \not\models \mathcal{O}'$ but Q_0 is equivalent to $(\mathcal{O}', \Sigma, q_0)$ as in fact both OMQs are equivalent to the CQ q_0 with atom $B(y)$ dropped.

The following is the main result of this section. When we speak about fixed-parameter tractability (FPT), we generally mean that the parameter is the size of the OMQ, that is, we refer to running time $f(\|Q\|) \cdot O(\|D\|^c)$ where f is a computable function and c a constant.

Theorem 1. *Let $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCT}\}$. Then \mathcal{ELT}_\perp^u -ontology relaxing OMQ evaluation is*

1. EXPTIME-complete in combined complexity and PTIME-complete in data complexity in $(\mathcal{L}, \mathcal{Q})$, $\mathcal{Q} \in \{AQ, CQ, UCQ\}$;
2. FPT with single exponential running time in $(\mathcal{L}, \mathcal{Q})$, $\mathcal{Q} \in \{CQ_k^{pw}, UCQ_k^{pw} \mid k \geq 1\}$;
3. in linear time in data complexity in $(\mathcal{L}, \text{bELIQ})$ with running time $2^{O(\|Q\|)} \cdot O(\|D\|)$.

To prove Theorem 1, we first establish the following crucial lemma that relates \mathcal{ELT}_\perp^u -ontology relaxing OMQ evaluation to tree unravelings.

Lemma 3. *Let $Q(\bar{x}) = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALCT}, \text{bELIQ})$, \mathcal{D} be a Σ -database, and $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$. Then*

1. $\bar{a} \in Q_{\mathcal{ELT}_\perp^u}(\mathcal{D})$ iff $\bar{a} \in Q(\mathcal{D}_\emptyset^\approx)$;
2. \mathcal{D} is satisfiable w.r.t. $\mathcal{O}_{\mathcal{ELT}_\perp^u}$ iff $\mathcal{D}_\emptyset^\approx$ is satisfiable w.r.t. \mathcal{O} .

We first prove Point 3 of Theorem 1. We are thus given an OMQ $Q(\bar{x}) = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALCT}, \text{bELIQ})$, a Σ -database \mathcal{D} , and a tuple $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$ and have to decide whether $\bar{a} \in Q_{\mathcal{ELT}_\perp^u}(\mathcal{D})$. By Point 1 of Lemma 3, it suffices to decide whether $\bar{a} \in Q(\mathcal{D}_\emptyset^\approx)$. This is much more convenient as we are back to the original ontology instead of having to deal directly with $\mathcal{O}_{\mathcal{ELT}_\perp^u}$. In the appendix, we show that deciding $\bar{a} \in Q(\mathcal{D}_\emptyset^\approx)$ can be reduced in linear time (data complexity) to the unsatisfiability of propositional Horn formulas, which is well-known to be in linear time (Dowling and Gallier 1984).

Regarding the upper bounds in Points 1 and 2, we first observe that we can concentrate on CQs rather than UCQs. This follows from the fact that \mathcal{ELT}_\perp^u -ontologies have universal models, even if infinite.

Lemma 4. *Let $Q(\bar{x}) = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALCT}, UCQ)$ with $q = p_1 \vee \dots \vee p_n$, and let \mathcal{D} be a Σ -database. Then $Q_{\mathcal{ELT}_\perp^u}(\mathcal{D}) = Q_1(\mathcal{D}) \cup \dots \cup Q_n(\mathcal{D})$, $Q_i = (\mathcal{O}_{\mathcal{ELT}_\perp^u}, \Sigma, p_i)$ for $1 \leq i \leq n$.*

We now describe an algorithm that establishes the upper bounds in Points 1 and 2 of Theorem 1 and that we refer to as a *careful chase*, see also (Bienvenu et al. 2013).

Assume that we are given an OMQ $Q(\bar{x}) = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALCT}, CQ)$, a Σ -database \mathcal{D} , and a tuple $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$. We use $\text{trees}(q)$ to denote the set of bELIQs that can be obtained from CQ q by first quantifying all variables, then taking a contraction, then an induced subquery, and then choosing at most one variable as the answer variable. In addition, $\text{trees}(q)$ contains all AQs $A(x)$ with A a concept name used in \mathcal{O} .

The algorithm first extends \mathcal{D} to a database \mathcal{D}' as follows:

- whenever $\mathcal{D}, \mathcal{O}_{\mathcal{ELT}_\perp^u} \models p()$ with $p() \in \text{trees}(q)$ a bELIQ, then take a disjoint copy of \mathcal{D}_p and add it to \mathcal{D} ;

- whenever $\mathcal{D}, \mathcal{O}_{\mathcal{ELT}_\perp^u} \models p(a)$ with $p(x) \in \text{trees}(q)$ an ELIQ, then take a disjoint copy of \mathcal{D}_p and add it to \mathcal{D} , glueing the root x of \mathcal{D}_p to a .

$\mathcal{D}, \mathcal{O}_{\mathcal{ELT}_\perp^u} \models p()$ and $\mathcal{D}, \mathcal{O}_{\mathcal{ELT}_\perp^u} \models p(a)$ can be decided in time $2^{O(\|Q\|)} \cdot O(\|D\|)$ by Point 3 of Theorem 1. Note that \mathcal{D}' is a subdatabase of $\text{ch}_{\mathcal{O}_{\mathcal{ELT}_\perp^u}}(\mathcal{D})$, which is why we speak of a careful chase. The algorithm then carries out the following steps:

- (i) If \mathcal{D} is unsatisfiable w.r.t. $\mathcal{O}_{\mathcal{ELT}_\perp^u}$, then returns ‘yes’;
- (ii) check whether $\bar{a} \in q(\mathcal{D}')$ and return the result.

By Point 2 of Lemma 3, the unsatisfiability check in (i) is equivalent to checking whether $\mathcal{D}_\emptyset^\approx$ is unsatisfiable w.r.t. \mathcal{O} . This is the case if and only if $\mathcal{D}_\emptyset^\approx, \mathcal{O} \models \exists x A(x)$ with A a fresh concept name, which can be decided in time $2^{O(\|Q\|)} \cdot O(\|D\|)$ by Point 3 of Theorem 1.

Checking $\bar{a} \in q(\mathcal{D}')$ in (ii) can be implemented using brute force to attain EXPTIME combined complexity and PTIME data complexity or using as a blackbox an algorithm that runs within the time requirements of fixed-parameter tractability to attain FPT when q is of bounded treewidth. We prove in the appendix that the algorithm is correct and achieves the upper bounds stated in Theorem 1.

The EXPTIME lower bounds in Theorem 1 are proved by a straightforward reduction from the subsumption of concept names in \mathcal{ALC} (Baader et al. 2017): A is subsumed by B w.r.t. \mathcal{O} iff $\{A(a)\}, \mathcal{O} \models B(a)$ iff $\{A(a)\}, \mathcal{O}_{\mathcal{ELT}_\perp^u} \models B(a)$. This trivial reduction also shows that using fragments of \mathcal{ELT}_\perp^u such as \mathcal{EL}_\perp or \mathcal{EL}_\perp^u as a target for ontology approximation cannot improve combined complexity. The PTIME lower bound in Point 1 is inherited from OMQ evaluation in (\mathcal{EL}, AQ) (Calvanese et al. 2015).

5 Tree-Database Relaxing Approximation

We study \mathcal{D}_\perp -database relaxing approximation that for the sake of readability we from now on refer to as tree-database relaxing approximation. We start with observing that tree-database relaxing approximation is incomparable to \mathcal{ELT}_\perp^u -ontology relaxing approximation.

Example 6. *Let $Q(x) = (\mathcal{O}, \Sigma, A(x)) \in (\mathcal{ALCT}, AQ)$ where*

$$\mathcal{O} = \{P \sqcap \exists r. P \sqsubseteq A, \neg P \sqcap \exists r. \neg P \sqsubseteq A\} \quad \Sigma = \{P, r, A\}.$$

Consider $\mathcal{D} = \{r(a, a)\}$. Then $a \notin Q_{\mathcal{ELT}_\perp^u}(\mathcal{D})$, but $a \in Q(\mathcal{D}_{\{a\}}^\approx)$ since $\mathcal{D}_{\{a\}}^\approx = \mathcal{D}$.

Conversely, let $Q(x) = (\emptyset, \{r\}, q) \in (\mathcal{ALCT}, CQ)$ where

$$q(x) = \exists y_1 \exists y_2 \exists y_3 r(x, y_1) \wedge r(y_1, y_2) \wedge r(y_2, y_3) \wedge r(y_3, y_1),$$

and $\mathcal{D} = \{r(a, b_1), r(b_1, b_2), r(b_2, b_3), r(b_3, b_1)\}$. Then $Q_{\mathcal{ELT}_\perp^u}(\mathcal{D}) = \{a\}$, but $Q(\mathcal{D}_{\{a\}}^\approx) = \emptyset$.

Note that the OMQs in Example 6 are based on CQs that are not bELIQs. This is no coincidence, as the following is a consequence of Lemma 3 and the fact that $\mathcal{D}_\emptyset^\approx \subseteq \mathcal{D}_S^\approx$ for all databases \mathcal{D} and $S \subseteq \text{adom}(\mathcal{D})$.

Proposition 1. In $(\mathcal{ALCI}, \text{bELIQ})$, tree-database relaxing OMQ evaluation is at least as complete as \mathcal{ELT}_\perp^u -ontology relaxing OMQ evaluation.

The converse of Proposition 1 fails, as per the first part of Example 6. The first main result of this section follows.

Theorem 2. Tree-database relaxing OMQ evaluation is

1. 2EXPTIME-complete in combined complexity and in linear time in data complexity (thus FPT) with running time $2^{2^{O(\|\mathcal{D}\|)}} \cdot O(\|\mathcal{D}\|)$ in $(\mathcal{ALCI}, \mathcal{Q})$, $\mathcal{Q} \in \{\text{CQ}, \text{UCQ}, \text{CQ}_k^{\text{nw}}, \text{UCQ}_k^{\text{nw}} \mid k \geq 1\}$;
2. EXPTIME-complete in combined complexity and in linear time in data complexity (thus FPT) with running time $2^{O(\|\mathcal{Q}\|)} \cdot O(\|\mathcal{D}\|)$ in $(\mathcal{ALC}, \mathcal{Q})$ and $(\mathcal{ALCI}, \mathcal{Q})$, $\mathcal{Q} \in \{\text{AQ}, \text{bELIQ}\}$.

The upper bound in Point 2 of Theorem 2 is proved similarly to Point 3 of Theorem 1, by reduction to the unsatisfiability of propositional Horn formulas. To prove the upper bound in Point 1, we first show that one can construct from an OMQ $Q = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALCI}, \text{UCQ})$ and Σ -database \mathcal{D} an OMQ $Q' = (\mathcal{O}, \Sigma', q')$ and a Σ' -database \mathcal{D}' such that $Q(\mathcal{D}) = Q'(\mathcal{D}')$ and q' is a conjunction of disjunctions of BELIQs. This again enables a reduction to the unsatisfiability of propositional Horn formulas. Note that a counterpart of Lemma 4 does not hold for tree-database relaxing approximation and thus we treat UCQs directly. The lower bounds are trivial as (non-approximate) evaluation in all mentioned OMQ languages is hard already on databases of the form $\mathcal{D} = \{A(a)\}$, which satisfy $\mathcal{D}_{\{a\}}^{\approx} = \mathcal{D}$, and for CQs of bounded treewidth (Lutz 2008a).

In contrast to \mathcal{ELT}_\perp^u -ontology relaxing approximation, we achieve no improvement in combined complexity in the \mathcal{ALCI} case, but we achieve linear time data complexity also for (U)CQs that are not of bounded treewidth. Informally, this is because database relaxing OMQ evaluation also approximates answers to the actual query (if it is a (U)CQ) while ontology relaxing evaluation only approximates the impact of the ontology. For comparison we recall that without ontologies, evaluating CQs of unbounded treewidth is $W[1]$ -hard, thus most likely not linear time in data complexity and not even FPT (Grohe 2007).

We next prove the second main result of this section. Recall that non-approximate OMQ evaluation in $(\mathcal{ALC}, \text{CQ})$ and $(\mathcal{ALC}, \text{UCQ})$ is EXPTIME-complete in combined complexity and CONP-complete in data complexity. Surprisingly, tree-database relaxing evaluation in these OMQ languages is *harder* in combined complexity than non-approximate evaluation while it is simpler in data complexity (the latter by Theorem 2).

Theorem 3. Tree-database relaxing OMQ evaluation is EXPSpace-hard in $(\mathcal{ALC}, \text{CQ})$ and 2EXPTIME-hard in $(\mathcal{ALC}, \text{UCQ})$.

The proof of Theorem 3 adapts hardness proofs from (Lutz 2008a) for non-approximate OMQ evaluation in $(\mathcal{ALCI}, \text{CQ})$, simulating inverse roles by making use of the unraveled database.

6 TGD-Ontology Relaxing Approximation

We generalize ontology relaxing approximation from \mathcal{ELT}_\perp^u to TGDs that are frontier-one and have bodies and heads of bounded treewidth. This yields an infinite hierarchy of increasingly more complete approximations.

A TGD $\vartheta = \phi(\bar{x}, \bar{y}) \rightarrow \exists \bar{z} \psi(\bar{x}, \bar{z})$ is an ℓ, k, ℓ', k' -TGD, $\ell < k$ and $\ell' < k'$, if its body has treewidth (ℓ, k) and its head has treewidth (ℓ', k') when viewed as CQs in which all variables are quantified. Every \mathcal{ELT}_\perp^u -CI can be written as a frontier-one 1, 2, 1, 2-TGD, but there are frontier-one 1, 2, 1, 2-TGDs that are not \mathcal{ELT}_\perp^u -CIs, such as $r(x, x) \rightarrow s(x, x)$. From now on, when speaking about ℓ, k, ℓ', k' -TGDs, we generally mean frontier-one TGDs. An ℓ, k, ℓ', k' -TGD-ontology is an FO-ontology that consists only of (frontier-one) ℓ, k, ℓ', k' -TGDs. If any of ℓ, k, ℓ', k' is not bounded, we assign to it value ω .

We study ℓ, k, ℓ', k' -TGD-ontology relaxing OMQ evaluation. Recall that, by Lemma 1, computing such approximations for an OMQ $Q \in (\text{FO}, \text{UCQ})$ corresponds to evaluating the OMQ $Q_{\ell, k, \ell', k'}^{\approx}$ which is based on ontology $\mathcal{O}_{\ell, k, \ell', k'}^{\approx}$. For brevity, we drop the suffix ‘-TGD’ and for instance speak about ℓ, k, ℓ', k' -ontology relaxing OMQ evaluation and write $Q_{\ell, k, \ell', k'}^{\approx}$.

We first observe that restricting the treewidth of the body is necessary to attain PTIME data complexity and that we cannot hope to attain the favourable combined complexity enjoyed by \mathcal{ELT}_\perp^u -ontology relaxing approximation for every value of ℓ' and k' . The following is a consequence of Example 7 below and of hardness proofs in (Lutz 2008a).

Proposition 2. $\omega, \omega, 1, 2$ -ontology relaxing OMQ evaluation in $(\mathcal{ALC}, \text{AQ})$ is CONP-hard in data complexity and there are ℓ', k' such that 1, 2, ℓ', k' -ontology relaxing OMQ evaluation in $(\mathcal{ALCI}, \text{CQ})$ is 2EXPTIME-hard in combined complexity.

To tame the combined complexity of TGD-ontology-relaxing approximation, we concentrate on the case where $\ell' = 1$. We next consider the choice of values for k and k' , the treewidth of rule bodies and heads. We first show that k gives rise to an infinite hierarchy of increasingly more complete approximations.

Example 7. Consider the OMQ $Q = (\mathcal{O}, \{e\}, \exists x D(x))$ from Example 2 that expresses non-3-colorability. For every $\{e\}$ -database \mathcal{D} , there is a $k \geq 1$ such that $\mathcal{D} \models Q$ iff $\mathcal{D} \models Q_{1, k, 1, 2}^{\approx}$: for $k = |\text{adom}(\mathcal{D})|$, $\mathcal{O}_{1, k, 1, 2}^{\approx}$ contains $q_{\mathcal{D}} \rightarrow \exists x D(x)$ if \mathcal{D} is not 3-colorable, where $q_{\mathcal{D}}$ is \mathcal{D} viewed as a CQ. Unless PTIME = NP, there are thus no ℓ, k, ℓ', k' such that $Q_{\ell, k, \ell', k'}^{\approx} \supseteq Q_{1, k+p, 1, 2}^{\approx}$ for all $p > 0$. In fact, the existence of such ℓ, k, ℓ', k' implies that 3-colorability is in PTIME by Theorem 4 below.

Example 7 also shows that TGD-ontology relaxing approximations are more complete than \mathcal{ELT}_\perp^u -ontology relaxing approximations, c.f. Example 2.

For k' , we make the weaker observation that there is no maximum value to be used uniformly for all OMQs.

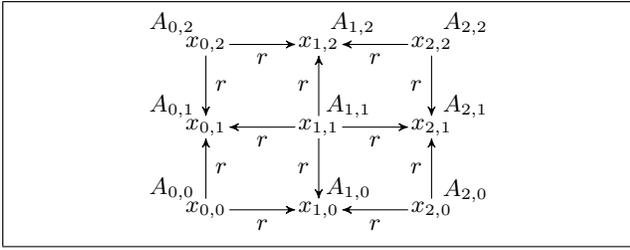


Figure 1: CQ from Example 8 for the case $m = 2$.

Example 8. For $m \geq 2$, consider the Boolean CQ

$$q_m = \bigwedge_{\substack{i,j,i',j' \in \{1,\dots,m\} \\ i+j \text{ is even} \\ |i-i'|+|j-j'|=1}} r(x_{i,j}, x_{i',j'}) \wedge \bigwedge_{i,j \in \{1,\dots,m\}} A_{i,j}(x_{i,j})$$

which takes the form of an $m \times m$ grid in which the direction of the roles alternate and every point in the grid is labeled with a unique concept name, see Figure 1. Consider the following two ways of contracting q_m into a path: identify $x_{i,j}$ with $x_{i',j'}$ if $i+j = i'+j'$ or identify $x_{i,j}$ with $x_{i',j'}$ if $i-j = i'-j'$. Let C_1 and C_2 be the two paths obtained, viewed as \mathcal{ELI} -concepts.

Set $\mathcal{O}_m = \{A \sqsubseteq C_1 \sqcup C_2\}$ and $Q_m() = (\mathcal{O}_m, \{A\}, q_m)$, and let $\mathcal{D} = \{A(a)\}$. Then $(\mathcal{O}_m)_{1,2,1,m^2}^{\approx}$ contains the TGD $A \rightarrow q_m$, so $\mathcal{D} \models (Q_m)_{1,2,1,m^2}^{\approx}$. In contrast, we argue in the appendix that $\mathcal{D} \not\models (Q_m)_{\omega,\omega,1,m^2-1}^{\approx}$.

For every fixed OMQ, however, there is a maximum useful value for k' . The next proposition is established analyzing the algorithm in the proof of Theorem 4 below.

Proposition 3. Let $\ell, k \geq 1$ with $\ell < k$. For every $Q(\bar{x}) = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALCI}, \text{UCQ})$ and $k' \geq |\text{var}(q)|$, $\mathcal{Q}_{\ell,k,1,|\text{var}(q)|}^{\approx} \equiv \mathcal{Q}_{\ell,k,1,k'}^{\approx}$.

The main result of this section is as follows.

Theorem 4. Let $\mathcal{L} \in \{\mathcal{ALC}, \mathcal{ALCI}\}$ and $\ell, k, k' \geq 1$ with $\ell < k$. Then $\ell, k, 1, k'$ -ontology relaxing OMQ evaluation is

1. EXPTIME-complete in combined complexity and PTIME-complete in data complexity in $(\mathcal{L}, \mathcal{Q})$, $\mathcal{Q} \in \{\mathcal{AQ}, \mathcal{CQ}, \text{UCQ}\}$;
2. FPT in $(\mathcal{L}, \mathcal{Q})$, $\mathcal{Q} \in \{\mathcal{CQ}_p^{hw}, \text{UCQ}_p^{hw} \mid p \geq 1\}$.

So TGD-ontology relaxing approximation inherits the good computational properties of \mathcal{ELI}_\perp^u -ontology relaxing approximation except for linear time for BELIQs, while being significantly more complete.

The lower bounds are proved exactly as for \mathcal{ELI}_\perp^u -ontology relaxing approximation, see Section 4. For the upper bounds, we treat the CQs in a UCQ independently and use a careful chase algorithm that essentially follows the lines of the careful chase presented in Section 4. An important difference is that a counterpart of Lemma 3 in which $Q_{\mathcal{ELI}_\perp^u}^{\approx}$ is replaced with $Q_{\ell,k,1,m^2}^{\approx}$ and $\mathcal{D}_\emptyset^{\approx}$ with an unraveling of \mathcal{D} into a database of treewidth ℓ, k fails to hold if $k' > k$ (the ‘only if’ direction of Point 1 fails). We resort to Theorem 5 below, which is a central ingredient to the proof. For

a database \mathcal{D} and $\ell, k \geq 1$, with $\ell < k$, $\mathcal{D}_{\ell,k}^{\approx}$ denotes the unraveling of \mathcal{D} into a database of treewidth (ℓ, k) , defined in the appendix. While $\text{adom}(\mathcal{D}) \cap \text{adom}(\mathcal{D}_{\ell,k}^{\approx}) = \emptyset$, $\mathcal{D}_{\ell,k}^{\approx}$ contains constants that are ‘copies’ of each $a \in \text{adom}(\mathcal{D})$. We use $\langle \mathcal{D}_{\ell,k}^{\approx}, a \rangle$ to denote a database obtained from $\mathcal{D}_{\ell,k}^{\approx}$ by choosing a copy of a in $\mathcal{D}_{\ell,k}^{\approx}$ and renaming it back to a . With $\langle \mathcal{D}_{\ell,k}^{\approx}, () \rangle$, we mean $\mathcal{D}_{\ell,k}^{\approx}$.

Theorem 5. Let $\ell, k, k' \geq 1$ with $\ell < k$. Given an OMQ $Q(\bar{x}) = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALCI}, \text{CQ})$ of arity at most one and with \mathcal{D}_q of treewidth $(1, k')$, a Σ -database \mathcal{D} , and $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$, deciding whether $\bar{a} \in Q(\langle \mathcal{D}_{\ell,k}^{\approx}, \bar{a} \rangle)$ is in EXPTIME in combined complexity and FPT.¹

The proof of Theorem 5 uses alternating tree automata.

Both the correctness proof for the careful chase algorithm and the proof of Theorem 5 exploit the regular shape of $\text{ch}_{\mathcal{O}}(\mathcal{D})$ for frontier one TGD-ontologies \mathcal{O} and fail for unrestricted TGDs. In fact, we do not know whether unrestricted TGD-ontology relaxing evaluation is decidable. However, we conjecture that the results in this section generalize to frontier-guarded TGDs (Baget, Leclère, and Mugnier 2010). The gain in completeness appears to be modest.

7 BTW-Database Relaxing Approximation

We study database relaxing approximation based on databases of bounded treewidth. For $\ell, k \geq 1$ with $\ell < k$, let $\mathcal{D}_{\ell,k}$ denote the class of pointed databases (\mathcal{D}, \bar{a}) such that the restriction of \mathcal{D} to domain $\text{adom}(\mathcal{D}) \setminus \bar{a}$ has treewidth ℓ, k . For readability, we speak of ℓ, k -database relaxing approximation in place of $\mathcal{D}_{\ell,k}$ -database relaxing approximation. As for TGD-ontology relaxing approximations, the parameter k gives rise to an infinite hierarchy of increasingly more complete approximations.

We first observe a counterpart of Lemma 2. Let \mathcal{D} be a database, \bar{a} a tuple over $\text{adom}(\mathcal{D})$, and $\ell, k \geq 1$ with $\ell < k$. With $\mathcal{D}_{\bar{a},\ell,k}^{\approx}$, we denote the database obtained by unraveling \mathcal{D} such that $(\mathcal{D}_{\bar{a},\ell,k}^{\approx}, \bar{a}) \in \mathcal{D}_{\ell,k}$. Details are in the appendix.

Lemma 5. Let $Q = (\mathcal{O}, \Sigma, q) \in (\text{FO}, \text{UCQ})$ and $\ell, k \geq 1$ with $\ell < k$. Then for all Σ -databases \mathcal{D} and $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$, $\bar{a} \in \text{app}_{\mathcal{D}_{\ell,k}}(Q, \mathcal{D})$ iff $\bar{a} \in Q(\mathcal{D}_{\bar{a},\ell,k}^{\approx})$.

We next relate ℓ, k, ℓ', k' -ontology relaxing approximation to ℓ, k -database relaxing approximation.

Example 9. Let $Q_n() = (\emptyset, \{r\}, q_n) \in (\mathcal{ALCI}, \text{CQ})$ where

$$q_n = \exists x_1 \cdots \exists x_n \bigwedge_{1 \leq i,j \leq n} r(x_i, x_j),$$

and let $\mathcal{D}_n = \{r(a_i, a_j) \mid 1 \leq i, j \leq n\}$. Then $\mathcal{D}_n \models (Q_n)_{1,2,1,2}^{\approx}$, but $(\mathcal{D}_n)_{\emptyset, n-2, n-1}^{\approx} \not\models Q_n$.

Conversely, take $Q() = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALCI}, \text{CQ})$ where

$$\begin{aligned} \mathcal{O} &= \{A \sqsubseteq \forall s. B \sqcup \forall s^2. B \sqcup \forall s^3. B\} \\ \Sigma &= \{A, A_1, A_2, A_3, B, r, s\} \\ q &= \exists x_1 \exists x_2 \exists x_3 \exists y \bigwedge_{1 \leq i \leq 3} r^2(x_i, y) \wedge A_i(x_i) \wedge B(y) \end{aligned}$$

¹Note that \mathcal{D}_q being of treewidth $(1, k')$ is a stricter condition than q being of treewidth $(1, k')$.

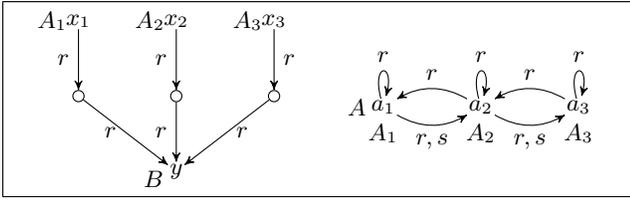


Figure 2: Query $q()$ and database \mathcal{D} from (2nd part of) Example 9.

and let

$$\mathcal{D} = \{A(a_1), A_i(a_i), s(a_i, a_{i+1}), r(a_i, a_{i+1}), r(a_{i+1}, a_i), r(a_i, a_i), r(a_3, a_3), A_3(a_3) \mid 1 \leq i < 3\}.$$

A picture is provided in Figure 2. The algorithms underlying the theorems in this and the previous section can be used to show that $\mathcal{D}_{\emptyset, 1, 2}^{\approx} \models Q$, but $\mathcal{D} \not\models Q_{\omega, \omega, \omega}^{\approx}$.

A straightforward variation of Example 7 shows that the parameter k indeed gives rise to an infinite hierarchy of increasingly more complete approximations.

Proposition 4. *Let $\ell, k, \ell', k' \geq 1$ with $\ell < k$ and $\ell' < k'$.*

1. In $(\mathcal{ALCI}, UCQ_{\ell, k}^{vw})$, ℓ, k -database relaxing OMQ evaluation is at least as complete as ℓ, k, ℓ', k' -ontology relaxing OMQ evaluation.
2. For OMQ from $(\mathcal{ALCI}, CQ_{\ell', k'}^{vw})$ of arity at most $r \leq 1$, $\ell + r, k + r, \ell', k'$ -ontology relaxing OMQ evaluation is at least as complete as ℓ, k -database relaxing OMQ evaluation.

We remark that Point 2 of Proposition 4 no longer holds if CQs are replaced by UCQs. We now formulate the main result of this section.

Theorem 6. *Let $\ell, k \geq 1$, $\ell < k$. Then ℓ, k -database relaxing OMQ evaluation is*

1. 2EXPTIME-complete in combined complexity and fixed-parameter tractable with double exponential running time in $(\mathcal{ALCI}, \mathcal{Q})$, $\mathcal{Q} \in \{CQ, UCQ, CQ_p^{lw}, UCQ_p^{lw} \mid p \geq 1\}$;
2. EXPTIME-complete in combined complexity and fixed-parameter tractable with single exponential running time in $(\mathcal{ALC}, \mathcal{Q})$ and $(\mathcal{ALCI}, \mathcal{Q})$, $\mathcal{Q} \in \{AQ, bELIQ\}$.

We thus achieve FPT even for (U)CQs of unbounded treewidth, as for tree-database relaxing approximations, but not linear time. The lower bounds are proved exactly as for Theorem 2. The upper bounds are shown by an elimination approach. We remark that ℓ, k -database relaxing approximation coincides with the answers given by canonical (ℓ, k) -Datalog programs, see (Feier, Kuusisto, and Lutz 2019).

Recall that for tree-database unraveling approximation, the combined complexity *increases* for OMQ languages based on \mathcal{ALC} and (U)CQs. It seems clear that the same is true for BTW-database relaxing approximation and that it can be proved by adapting the proof of Theorem 3. A central idea is to replace single constants in databases by cliques.

Conjecture 1. *Let $\ell, k \geq 1$, $\ell < k$. Then ℓ, k -database relaxing OMQ evaluation in (\mathcal{ALC}, UCQ) is 2EXPTIME-hard.*

8 Approximation from Above

The approximations studied so far are from below, thus sound but incomplete. We define dual approximations from above that are complete but unsound: ontology strengthening approximation and database strengthening approximation. It turns out that these are computationally less well-behaved. While the former may increase combined complexity, the latter does not even enjoy PTIME data complexity.

We start with ontology strengthening approximation. For an OMQ $Q(\bar{x}) = (\mathcal{O}, \Sigma, q)$, a Σ -database \mathcal{D} , and an ontology language \mathcal{L}' , we use $\text{app}_{\mathcal{L}'}^{\uparrow}(Q, \mathcal{D})$ to denote the set of tuples $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$ such that $\bar{a} \in Q'(\mathcal{D})$ for all OMQs $Q' = (\mathcal{O}', \Sigma, q)$ where \mathcal{O}' is an \mathcal{L}' -ontology with $\mathcal{O}' \models \mathcal{O}$. Every choice of $(\mathcal{L}, \mathcal{Q})$ and \mathcal{L}' gives rise to an approximate OMQ evaluation problem.

\mathcal{L}' -ontology strengthening evaluation in $(\mathcal{L}, \mathcal{Q})$
 INPUT : OMQ $Q(\bar{x}) = (\mathcal{O}, \Sigma, q) \in (\mathcal{L}, \mathcal{Q})$,
 Σ -database \mathcal{D} , tuple $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$
 OUTPUT : ‘yes’ if $\bar{a} \in \text{app}_{\mathcal{L}'}^{\uparrow}(Q, \mathcal{D})$ and ‘no’ otherwise

We consider \mathcal{ELI}_{\perp} -ontology strengthening OMQ evaluation in $(\mathcal{ELIU}_{\perp}, UCQ)$ where \mathcal{ELIU}_{\perp} is the extension of \mathcal{ELI}_{\perp} with disjunction. Note that we can find an implying \mathcal{ELI}_{\perp} -ontology for every \mathcal{ELIU}_{\perp} -ontology, namely $\{\top \sqsubseteq \perp\}$. In contrast to ontology relaxing approximation, it does not seem beneficial to use the universal role. The following example illustrates unsoundness.

Example 10. *Take $Q(x) = (\mathcal{O}, \Sigma, q)$ where*

$$\begin{aligned} \mathcal{O} &= \{A \sqsubseteq A_1 \sqcup A_2, \\ &\quad \exists r.(A_i \sqcap B_1) \sqcap \exists r.(A_i \sqcap B_2) \sqsubseteq B \mid i \in \{1, 2\}\} \\ \Sigma &= \{A, A_1, A_2, B, B_1, B_2, r\} \\ q &= B(x). \end{aligned}$$

Let $\mathcal{D} = \{r(a, b_1), r(a, b_2), A(b_1), B_1(b_1), A(b_2), B_2(b_2)\}$. Then $a \in \text{app}_{\mathcal{ELI}_{\perp}}^{\uparrow}(\mathcal{D})$ as every \mathcal{ELI}_{\perp} -ontology \mathcal{O}' with $\mathcal{O}' \models \mathcal{O}$ implies $A \sqsubseteq A_1$ or $A \sqsubseteq A_2$. But $a \notin Q(\mathcal{D})$.

Let \mathcal{O} be an \mathcal{ELIU}_{\perp} -ontology. A set \mathcal{M} of \mathcal{ELI}_{\perp} -ontologies is an *exhaustive \mathcal{ELI}_{\perp} -approximation set* for \mathcal{O} if $\hat{\mathcal{O}} \models \mathcal{O}$ for every $\hat{\mathcal{O}} \in \mathcal{M}$, and for every \mathcal{ELI}_{\perp} -ontology \mathcal{O}' with $\mathcal{O}' \models \mathcal{O}$, there is an $\hat{\mathcal{O}} \in \mathcal{M}$ such that $\mathcal{O}' \models \hat{\mathcal{O}}$. Such sets \mathcal{M} are interesting because for all OMQs $Q = (\mathcal{O}, \Sigma, q)$ with q a UCQ and all Σ -databases \mathcal{D} , $\text{app}_{\mathcal{ELI}_{\perp}}^{\uparrow}(Q, \mathcal{D}) = \bigcap_{\hat{\mathcal{O}} \in \mathcal{M}} Q_{\hat{\mathcal{O}}}(\mathcal{D})$ where $Q_{\hat{\mathcal{O}}} = (\hat{\mathcal{O}}, \Sigma, q)$. Consider for instance $\mathcal{O} = \{\top \sqsubseteq A_1 \sqcup A_2\}$, for which $\{\mathcal{O}_1, \mathcal{O}_2\}$ is an exhaustive \mathcal{ELI}_{\perp} -approximation set where $\mathcal{O}_i = \{\top \sqsubseteq A_i\}$. In the appendix, we show how to construct a finite exhaustive \mathcal{ELI}_{\perp} -approximation set for any given \mathcal{ELIU}_{\perp} -ontology and use this to prove the upper bounds in the following result.

Theorem 7. *Let $\mathcal{L} \in \{AQ, CQ, UCQ\}$. \mathcal{ELI}_{\perp} -ontology strengthening OMQ evaluation in $(\mathcal{ELIU}_{\perp}, \mathcal{L})$ is 2EXPTIME-complete in combined complexity and FPT with double exponential running time.*

	Ontology relaxing		Database relaxing	
	\mathcal{ELI}_\perp^u	TGD	trees	bounded treewidth
$(\mathcal{ALCC}(\mathcal{I}), \text{bELIQ})$	LIN / EXP TIME	FPT / EXP TIME	LIN / EXP TIME	FPT / EXP TIME
$(\mathcal{ALCC}, \text{CQ})$	P TIME / EXP TIME		FPT / EXP SPACE-2 EXP TIME	FPT / in 2 EXP TIME
$(\mathcal{ALCC}, \text{UCQ})$	P TIME / EXP TIME		FPT / 2 EXP TIME	FPT / in 2 EXP TIME
$(\mathcal{ALCCI}, (\text{UCQ})_k^{\text{nw}})$	FPT / EXP TIME		FPT / 2 EXP TIME	
$(\mathcal{ALCCI}, (\text{UCQ}))$	P TIME / EXP TIME		FPT / 2 EXP TIME	

Figure 3: Results for approximate OMQ evaluation, data and parametric complexity / combined complexity.

The lower bound is proved by a reduction from the word problem for a suitable kind of alternating Turing machine. We consider the lower bound for $(\mathcal{ELIU}_\perp, \text{AQ})$ surprising as non-approximate OMQ evaluation is only EXP TIME-complete (Lutz 2008a). Thus, approximate OMQ evaluation is significantly harder, the only result of this kind in the current paper that applies to AQs. It is amusing to note that the lower bound depends only on disjunction on the left hand side of concept inclusions, which are syntactic sugar, but not on the seemingly much more ‘dangerous’ disjunctions on the right hand side. It is in fact a byproduct of our proofs that, without disjunctions on the left, \mathcal{ELI}_\perp -ontology strengthening OMQ evaluation in $(\mathcal{ELIU}_\perp, \text{UCQ})$ is only EXP TIME-complete. \mathcal{ALCCI} -ontologies can be rewritten in polynomial time into a ‘nesting-free’ normal form that is often used by reasoners and that has sometimes been presupposed for approximation (Zhou et al. 2015). The rewriting is not equivalence preserving, but only yields a conservative extension. \mathcal{ALCCI} -ontologies in this form can in turn be rewritten into an equivalent \mathcal{ELIU}_\perp -ontology without disjunction on the left. The following example shows that \mathcal{ALCC} -ontologies that are not in normal form behave differently in that they may have only infinite exhaustive \mathcal{ELI}_\perp -approximation sets.

Example 11. Let $\mathcal{O} = \{\exists r. \top \sqcap \forall r. A \sqsubseteq B_1 \sqcup B_2\}$. Then for each $n \geq 1$, the \mathcal{ELI}_\perp -ontology

$$\mathcal{O}_n = \{\exists r. A \sqsubseteq \exists r^n. X, \exists r. (A \sqcap \exists r^{n-1}. X) \sqsubseteq B_1\}$$

is such that $\mathcal{O}_n \models \mathcal{O}$. It is easy to see that $\mathcal{O}_n \not\models \mathcal{O}_m$ when $n \neq m$ and any \mathcal{ELI}_\perp -ontology \mathcal{O}'_n with $\mathcal{O}_n \models \mathcal{O}'_n \models \mathcal{O}$ is equivalent to \mathcal{O}_n .

In the appendix, we give another example which shows that the effect pointed out in Example 11 also affects answers to OMQs. We leave the decidability and complexity of \mathcal{ELI}_\perp -ontology relaxing approximation in $(\mathcal{ALCCI}, \text{UCQ})$ (without assuming normal form) as an open problem.

We next turn to database strengthening approximation. For an OMQ $Q(\bar{x}) = (\mathcal{O}, \Sigma, q)$, a Σ -database \mathcal{D} , and a class \mathfrak{D} of pointed databases, we use $\text{app}_{\mathfrak{D}}^\uparrow(Q, \mathcal{D})$ to denote the set of tuples $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$ such that for all $\langle \mathcal{D}', \bar{b} \rangle \in \mathfrak{D}$ and all homomorphisms h from \mathcal{D} to \mathcal{D}' with $h(\bar{a}) = \bar{b}$, $\bar{b} \in Q(\mathcal{D}')$. Every choice of $(\mathcal{L}, \mathcal{Q})$ and \mathfrak{D} gives rise to an approximate OMQ evaluation problem.

\mathfrak{D} -database strengthening evaluation in $(\mathcal{L}, \mathcal{Q})$
INPUT : OMQ $Q(\bar{x}) = (\mathcal{O}, \Sigma, q) \in (\mathcal{L}, \mathcal{Q})$,
 Σ -database \mathcal{D} , tuple $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$
OUTPUT : ‘yes’ if $\bar{a} \in \text{app}_{\mathfrak{D}}^\uparrow(Q, \mathcal{D})$ and ‘no’ otherwise

A natural choice for \mathfrak{D} are classes of databases of bounded treewidth. We only consider here the class \mathfrak{D}_1 of pointed databases $\langle \mathcal{D}, \bar{a} \rangle$ where $\text{adom}(\mathcal{D}) \setminus \bar{a}$ has treewidth $(1, 2)$, i.e. it is a tree with multi-edge and self-loops admitted.

Example 12. Take $Q(x) = (\mathcal{O}, \Sigma, q)$ where

$$\mathcal{O} = \{A_i \sqcap A_j \sqsubseteq B \mid 1 \leq i < j \leq 3\}$$

$$\Sigma = \{A_1, A_2, A_3, B, r\}$$

$$q = \exists x B(x).$$

Let

$$\mathcal{D} = \{r(a_1, a_2), r(a_2, a_3), r(a_3, a_1), A_1(a_1), A_2(a_2), A_3(a_3)\}.$$

Then $() \in \text{app}_{\mathfrak{D}_1}^\uparrow(\mathcal{D})$, but $() \notin Q(\mathcal{D})$.

Examples 10 and 12 also show that \mathfrak{D}_1 -database strengthening approximation and \mathcal{ELI}_\perp -ontology strengthening approximation are incomparable as $a \notin \text{app}_{\mathfrak{D}_1}^\uparrow(\mathcal{D})$ in Example 10 and $() \notin \text{app}_{\mathcal{ELI}_\perp}^\uparrow(\mathcal{D})$ in Example 12. It turns out that \mathfrak{D}_1 -database strengthening approximation does not even enjoy P TIME data complexity, and this holds already when the original ontology is formulated in a tractable language.

Theorem 8. \mathfrak{D}_1 -database strengthening approximation is CONP-complete in data complexity in $(\mathcal{ALCCI}, \text{UCQ})$. The lower bound already holds when the ontology is empty. It also holds in $(\mathcal{EL}, \text{CQ})$.

The interesting part is the lower bound, proved by a non-trivial reduction from the validity of propositional formulas.

9 Conclusion

We have introduced and studied various kinds of OMQ approximations. Our results on approximation from below are summarized in Figure 3 where all entries for combined complexity mean completeness results, unless stated otherwise. We believe that the results in this paper show that ontology relaxing approximation is preferable to database-relaxing approximation in most aspects. First, they are conservative in the sense that they are complete when the ontology is empty. Second, they have lower combined complexity. And third, also for the dual notion of ontology-strengthening approximation one attains P TIME data complexity. It would be interesting to generalize our approach to extensions of \mathcal{ALCCI} with, for instance, role hierarchies, transitive roles, and functional roles. Moreover, ontology strengthening approximation deserves further study.

Acknowledgements

Haga and Lutz were supported by DFG CRC 1320 Ease. Wolter was supported by EPSRC grant EP/S032207/1.

References

- Baader, F.; Horrocks, I.; Lutz, C.; and Sattler, U. 2017. *An Introduction to Description Logic*. Cambridge University Press.
- Baget, J.; Leclère, M.; Mugnier, M.; and Salvat, E. 2009. Extending decidable cases for rules with existential variables. In *Proc. of IJCAI*, 677–682.
- Baget, J.; Leclère, M.; and Mugnier, M. 2010. Walking the decidability line for rules with existential variables. In *Proc. of KR*. AAAI Press.
- Barceló, P.; Feier, C.; Lutz, C.; and Pieris, A. 2019. When is ontology-mediated querying efficient? In *Proc. of LICS*, 1–13. IEEE.
- Barceló, P.; Libkin, L.; and Romero, M. 2014. Efficient approximations of conjunctive queries. *SIAM J. Comput.* 43(3):1085–1130.
- Barceló, P.; Romero, M.; and Zeume, T. 2018. A more general theory of static approximations for conjunctive queries. In *Proc. of ICDT*, volume 98 of *LIPICs*, 7:1–7:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Bienvenu, M.; Ortiz, M.; Simkus, M.; and Xiao, G. 2013. Tractable queries for lightweight description logics. In *Proc. of IJCAI*, 768–774. IJCAI/AAAI.
- Bienvenu, M.; ten Cate, B.; Lutz, C.; and Wolter, F. 2014. Ontology-based data access: A study through disjunctive datalog, CSP, and MMSNP. *ACM Trans. Database Syst.* 39(4):33:1–33:44.
- Bötcher, A.; Lutz, C.; and Wolter, F. 2019. Ontology approximation in Horn description logics. In *Proc. of IJCAI*, 1574–1580. ijcai.org.
- Calì, A.; Gottlob, G.; and Kifer, M. 2013. Taming the infinite chase: Query answering under expressive relational constraints. *J. Artif. Intell. Res.* 48:115–174.
- Calvanese, D.; Giacomo, G. D.; Lembo, D.; Lenzerini, M.; and Rosati, R. 2015. Data complexity of query answering in description logics (extended abstract). In *Proc. of IJCAI*, 4163–4167. AAAI Press.
- Carral, D.; Feier, C.; Grau, B. C.; Hitzler, P.; and Horrocks, I. 2014. \mathcal{EL} -ifying ontologies. In *Proc. of IJCAR*, volume 8562 of *LNCS*, 464–479. Springer.
- Chandra, A. K.; Kozen, D.; and Stockmeyer, L. J. 1981. Alternation. *J. ACM* 28(1):114–133.
- Dechter, R., and Pearl, J. 1989. Tree clustering for constraint networks. *Artif. Intell.* 38(3):353–366.
- Dowling, W. F., and Gallier, J. H. 1984. Linear-time algorithms for testing the satisfiability of propositional Horn formulae. *J. Log. Program.* 1(3):267–284.
- Eiter, T.; Gottlob, G.; Ortiz, M.; and Simkus, M. 2008. Query answering in the description logic Horn-*SHIQ*. In *Proc. of JELIA*, volume 5293 of *LNCS*, 166–179. Springer.
- Fan, W., and Geerts, F. 2012. *Foundations of Data Quality Management*. Synthesis Lectures on Data Management. Morgan & Claypool Publishers.
- Feier, C.; Kuusisto, A.; and Lutz, C. 2019. Rewritability in monadic disjunctive Datalog, MMSNP, and expressive description logics. *Log. Methods Comput. Sci.* 15(2).
- Fink, R., and Olteanu, D. 2011. On the optimal approximation of queries using tractable propositional languages. In *Proc. of ICDT*, 174–185. ACM.
- Freuder, E. C. 1990. Complexity of k-tree structured constraint satisfaction problems. In *Proc. of AAAI*, 4–9. AAAI Press / The MIT Press.
- Gottlob, G.; Morak, M.; and Pieris, A. 2015. Recent advances in datalog[±]. In *Reasoning Web*, volume 9203 of *LNCS*, 193–217. Springer.
- Grohe, M. 2007. The complexity of homomorphism and constraint satisfaction problems seen from the other side. *J. ACM* 54(1):1:1–1:24.
- Haga, A.; Lutz, C.; Marti, J.; and Wolter, F. 2020. A journey into ontology approximation: From non-Horn to Horn. In *Proc. of IJCAI*, 1827–1833. ijcai.org.
- Haga, A.; Lutz, C.; Sabellek, L.; and Wolter, F. 2021. How to approximate ontology-mediated queries. *CoRR* abs/2107.05369.
- Hernich, A.; Lutz, C.; Papacchini, F.; and Wolter, F. 2020. Dichotomies in ontology-mediated querying with the guarded fragment. *ACM Trans. Comput. Log.* 21(3):20:1–20:47.
- Hitzler, P., and Vrandečić, D. 2005. Resolution-based approximate reasoning for OWL DL. In *Proc. of ISWC*, 383–397. Springer.
- Johnson, D. S., and Klug, A. C. 1982. Testing containment of conjunctive queries under functional and inclusion dependencies. In *Proc. of PODS*, 164–169. ACM.
- Kaminski, M.; Nenov, Y.; and Grau, B. C. 2016. Datalog rewritability of disjunctive datalog programs and non-Horn ontologies. *Artif. Intell.* 236:90–118.
- Lutz, C., and Wolter, F. 2017. The data complexity of description logic ontologies. *Log. Methods Comput. Sci.* 13(4).
- Lutz, C. 2008a. The complexity of conjunctive query answering in expressive description logics. In *Proc. of IJCAR*, volume 5195 of *LNCS*, 179–193. Springer.
- Lutz, C. 2008b. Two upper bounds for conjunctive query answering in *SHIQ*. In *Proc. of DL2008*, volume 353 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Pan, J. Z., and Thomas, E. 2007. Approximating OWL-DL ontologies. In *Proc. of AAAI*, 1434–1439. AAAI Press.
- Schaerf, A. 1993. On the complexity of the instance checking problem in concept languages with existential quantification. *J. of Intel. Inf. Systems* 2:265–278.
- Selman, B., and Kautz, H. A. 1996. Knowledge compilation and theory approximation. *J. ACM* 43(2):193–224.
- Thomas, E.; Pan, J. Z.; and Ren, Y. 2010. TrOWL: Tractable OWL 2 reasoning infrastructure. In *Proc. of ESWC*, volume 6089 of *LNCS*, 431–435. Springer.
- Tserendorj, T.; Rudolph, S.; Krötzsch, M.; and Hitzler, P. 2008. Approximate OWL-reasoning with Screech. In *Proc. of RR*, volume 5341 of *LNCS*, 165–180. Springer.

Vardi, M. Y. 1985. The taming of converse: Reasoning about two-way computations. In *Logics of Programs*, volume 193 of *LNCS*, 413–423. Springer.

Zhou, Y.; Cuenca Grau, B.; Nenov, Y.; Kaminski, M.; and Horrocks, I. 2015. PAGOdA: Pay-as-you-go ontology query answering using a datalog reasoner. *J. Artif. Intell. Res.* 54:309–367.

A Further Preliminaries

A.1 The Chase

Let $\vartheta = \phi(\bar{x}, \bar{y}) \rightarrow \exists \bar{z} \psi(\bar{x}, \bar{z})$ be a TGD and \mathcal{D} a database. If h is a homomorphism from $\phi(\bar{x}, \bar{y})$ to \mathcal{D} , then we say that ϑ is *applicable to \mathcal{D} via h* . Now assume that ϑ is applicable to \mathcal{D} via h . Let $h'(\psi(\bar{x}, \bar{z}))$ be obtained from $\psi(\bar{x}, \bar{z})$ by replacing the variables x in \bar{x} by $h(x)$ and replacing the variables z in \bar{z} with fresh constants. Then *the result of applying ϑ to \mathcal{D} via h* is defined as the database $\mathcal{D}' = \mathcal{D} \cup h'(\psi(\bar{x}, \bar{z}))$ and we write $\mathcal{D} \xrightarrow{h, \vartheta} \mathcal{D}'$.

Let \mathcal{D} be a database and \mathcal{O} be a set of TGDs, both potentially infinite, but at most countable. A *chase sequence of \mathcal{D} w.r.t. \mathcal{O}* is a potentially infinite sequence $\mathcal{D}_0, \mathcal{D}_1, \dots$ such that $\mathcal{D}_0 = \mathcal{D}$ and for every $i > 0$, there are a homomorphism h_i and a TGD $\vartheta_i \in \mathcal{O}$ with $\mathcal{D}_{i-1} \xrightarrow{h_i, \vartheta_i} \mathcal{D}_i$. The chase sequence $\mathcal{D}_0, \mathcal{D}_1, \dots$ is *fair* if for every $i \geq 0$ and $\vartheta \in \mathcal{O}$ that is applicable to \mathcal{D}_i via some homomorphism h , there is a $j > i$ such that \mathcal{D}_{j+1} is the result of applying ϑ to \mathcal{D}_j via h . We then denote by $\text{ch}_{\mathcal{O}}(\mathcal{D})$ the union of all \mathcal{D}_i for some fair chase sequence of \mathcal{D} w.r.t. \mathcal{O} and call it the *result of chasing \mathcal{D} with \mathcal{O}* . Note that our chase is oblivious, that is, it may apply a TGD via a homomorphism h despite the fact that h can be extended to a homomorphism from the head to the database. As a consequence, the result of the chase is unique up to isomorphism.

Lemma 6. *Let \mathcal{O} be a (potentially infinite) TGD-ontology and \mathcal{D} a (potentially infinite) Σ -database that is satisfiable w.r.t. \mathcal{O} . Then the following holds:*

1. $\text{ch}_{\mathcal{O}}(\mathcal{D})$ is a model of \mathcal{D} and \mathcal{O} ;
2. for every model \mathcal{I} of \mathcal{D} and \mathcal{O} there exists a homomorphism h from $\text{ch}_{\mathcal{O}}(\mathcal{D})$ to \mathcal{I} with $h(a) = a$ for all $a \in \text{adom}(\mathcal{D})$;
3. for $Q = (\mathcal{O}, \Sigma, q)$ with q a UCQ and all $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$, $\bar{a} \in Q(\mathcal{D})$ iff $\bar{a} \in q(\text{ch}_{\mathcal{O}}(\mathcal{D}))$.

Proof. We provide a sketch only. Condition 1 holds by definition of $\text{ch}_{\mathcal{O}}(\mathcal{D})$. For Condition 2, let the chase sequence that constructs $\text{ch}_{\mathcal{O}}(\mathcal{D})$ be $\mathcal{D}_0, \mathcal{D}_1, \dots$ and let \mathcal{I} be a model of \mathcal{D} and \mathcal{O} . One can construct homomorphisms h_i from \mathcal{D}_i to \mathcal{I} , for all $i \geq 0$, by induction on i and starting with the identity mapping h_0 from \mathcal{D} to \mathcal{I} in a straightforward way. Note that it does not play a role here that \mathcal{O} might be infinite. The only relevant point is that we use countable sets of symbols, which guarantees that \mathcal{O} is countable as well and thus a fair chase sequence always exists. Condition 3 is a consequence of 1 and 2. \square

Since every \mathcal{ELI}_{\perp} -ontology can be viewed as a TGD-ontology (see Section 2), we may also apply the chase directly to \mathcal{ELI}_{\perp} -ontologies.

A.2 Unraveling Databases

We define the unraveling of a database \mathcal{D} into a database of bounded treewidth. Let $1 \leq \ell < k$, and $S \subseteq \text{adom}(\mathcal{D})$. The ℓ, k -unraveling $\mathcal{D}_{S, \ell, k}^{\approx}$ of \mathcal{D} up to S is defined as follows. An ℓ, k -sequence takes the form

$$v = S_0, O_0, S_1, O_1, S_2, \dots, O_{n-1}, S_n,$$

$n \geq 0$, where $S \subseteq S_i \subseteq \text{adom}(\mathcal{D})$, $O_i \subseteq \text{adom}(\mathcal{D})$, and the following conditions are satisfied:

- $|S_i \setminus S| \leq k$ for $0 \leq i \leq n$;
- $S \subseteq O_i \subseteq S_i \cap S_{i+1}$ for $0 \leq i < n$;
- $|O_i \setminus S| \leq \ell$ for $0 \leq i < n$.

For every $a \in \text{adom}(\mathcal{D})$, reserve a countably infinite set of fresh constants that we refer to as *copies* of a . For brevity, we also consider a to be a copy of itself. A *bag* is a finite set of copies of constants in $\text{adom}(\mathcal{D})$.

Now let (V, E) be the infinite directed tree with V the set of all ℓ, k -sequences and E the prefix order on V . We proceed inductively on n to choose, for every $v = S_0 \dots S_n \in V$, a bag B_v that contains a copy a_v of every $a \in S_n$ and no other constants such that

1. if $a \in S$, then the copy a_v of a is a itself;
2. if $n > 0$ and $a \in O_{n-1}S$, then the copy a_v of a is identical to the copy $a_{S_0 \dots S_{n-1}}$ of a ;
3. if $n = 0$ and $a \notin S$ or $n > 0$ and $a \notin O_{n-1}$, then the copy a_v of a is fresh.

We then define $\mathcal{D}_{S, \ell, k}^{\approx}$ as the (unique and infinite) database with active domain $\bigcup_{v \in V} B_v$ such that

1. $(V, E, (B_v)_{v \in V})$ is a tree decomposition of $\mathcal{D}_{S, \ell, k}^{\approx}$ and
2. the ‘uncopying’ map $a_v \mapsto a$ is an isomorphism from $(\mathcal{D}_{S, \ell, k}^{\approx})|_{B_v}$ onto $\mathcal{D}|_{S_n}$, for every $v = S_0 \dots S_n \in V$.

It is easy to see that the ‘uncopying’ map is a homomorphism from $\mathcal{D}_{S, \ell, k}^{\approx}$ onto \mathcal{D} . If $S = \emptyset$, then we drop S and simply speak of the ℓ, k -unraveling $\mathcal{D}_{\ell, k}^{\approx}$ of \mathcal{D} .

We next give some properties of unraveled databases that are fundamental to the remainder of this paper. We first observe that homomorphisms from databases of treewidth (ℓ, k) into a database \mathcal{D} can be lifted to homomorphisms into the ℓ, k -unraveling of \mathcal{D} .

Lemma 7. *Let (\mathcal{D}, \bar{a}) be a pointed database and $1 \leq \ell < k$.*

1. *Let (\mathcal{D}', \bar{b}) be a pointed database of treewidth (ℓ, k) . If there is a homomorphism h from \mathcal{D}' to \mathcal{D} with $h(\bar{b}) = \bar{a}$, then there is a homomorphism g from \mathcal{D}' to $\mathcal{D}_{\bar{a}, \ell, k}^{\approx}$ with $g(\bar{b}) = \bar{a}$.*
2. *Let \mathcal{D}' be a database of treewidth (ℓ, k) . If there is a homomorphism h from \mathcal{D}' to \mathcal{D} with $h(b) = a$, then for any copy a' of a in $\mathcal{D}_{\bar{a}, \ell, k}^{\approx}$ there is a homomorphism g from \mathcal{D}' to $\mathcal{D}_{\bar{a}, \ell, k}^{\approx}$ with $g(b) = a'$.*

Proof. (1) Assume that (\mathcal{D}', \bar{b}) is a pointed database of treewidth (ℓ, k) . Then the restriction of \mathcal{D}' to $\text{adom}(\mathcal{D}') \setminus \bar{b}$ has treewidth (ℓ, k) . Assume there is a homomorphism h from \mathcal{D}' to \mathcal{D} such that $h(\bar{b}) = \bar{a}$. To define g , let $(V, E, (B_v)_{v \in V})$ be an (ℓ, k) -tree decomposition of the restriction of \mathcal{D}' to $\text{adom}(\mathcal{D}') \setminus \bar{b}$. We may assume that (V, E) is connected and start defining g by setting $g(\bar{b}) = h(\bar{b}) = \bar{a}$. Now pick a first $v \in V$. We aim to define $g(b)$ for $b \in B_v$. Let (V', E') be the underlying tree of the tree decomposition of $\mathcal{D}_{\bar{a}, \ell, k}^{\approx}$ introduced in the definition of $\mathcal{D}_{\bar{a}, \ell, k}^{\approx}$ and recall that V' consists of ℓ, k -sequences. To guide the construction, we

also define a homomorphism g' from (V, E) and (V', E') and start by setting $g'(v) = h(B_v) \cup \bar{a}$. Define $g(b)$ as the copy of $h(b)$ in $B_{g'(v)}$, for all $b \in B_v$. Next assume that g' has been defined on $U \subseteq V$ and g on $\bigcup_{v \in U} B_v$ with U connected and that $(v, v') \in E$ with $v \in U$ but $v' \notin U$. Then $g'(v)$ is an ℓ, k -sequence and we can expand g' by setting $g'(v') = g'(v)(h(B_v \cap B_{v'}) \cup \bar{a})(h(B_{v'}) \cup \bar{a})$. The definition of g on $B_{v'}$ is as expected by setting $g(b) = h(b)'$ for the copy $h(b)'$ of $h(b)$ in $B_{g'(v')}$. Then g restricted to B_v is clearly a homomorphism. Overall, it follows that the restriction of g to $\text{adom}(\mathcal{D}') \setminus \bar{b}$ is a homomorphism to $\mathcal{D}_{\bar{a}, \ell, k}^{\approx}$. That g is a homomorphism from \mathcal{D}' to $\mathcal{D}_{\bar{a}, \ell, k}^{\approx}$ is now a consequence of the fact that h is homomorphism and Point 2 of the definition of $\mathcal{D}_{\bar{a}, \ell, k}^{\approx}$.

(2) The proof is similar to the proof Claim (1). In this case we have to start the definition of g and g' from some fixed $g(b) = a'$ with a' a copy of $a = h(b)$. We use the notation from (1) and assume that $b \in B_v$. Consider a shortest ℓ, k -sequence $w = S_0 \cdots S_n$ such that $a' \in B_w$. Then we set $g'(v) = wO_n(h(B_v) \cup \bar{a})$, where $O_n = \{a'\} \cup \bar{a}$. Define $g(b')$ as the copy of $h(b')$ in $B_{g'(v)}$, for all $b' \in B_v$. Observe that then $g(b) = a'$, as required. The function g can now be extended to a homomorphism from \mathcal{D}' to $\mathcal{D}_{\bar{a}, \ell, k}^{\approx}$ in exactly the same way as above. \square

The subsequent two lemmas are related to the notion of unraveling tolerance considered in (Lutz and Wolter 2017) where it is observed that when a database \mathcal{D} is unraveled into a proper tree $\mathcal{D}_{\emptyset}^{\approx}$ and $a \in \text{adom}(\mathcal{D})$, then for any OMQ Q from $(\mathcal{ELI}_{\perp}, \text{ELIQ})$, a is an answer to Q on \mathcal{D} iff it is an answer to Q on $\mathcal{D}_{\emptyset}^{\approx}$, see Lemma 10 below. In our case, ontologies are sets of TGDs, and unraveling only provides databases of bounded treewidth. This makes the formulation more subtle as we have to make sure that all parameters fit together.

Lemma 8. *Let \mathcal{D} be a database, $S \subseteq \text{adom}(\mathcal{D})$, and \mathcal{O} a possibly infinite ℓ, k, ℓ', k' -ontology, $1 \leq \ell < k$ and $1 \leq \ell' < k'$. Then \mathcal{D} is satisfiable w.r.t. \mathcal{O} iff $\mathcal{D}_{S, \ell, k}^{\approx}$ is satisfiable w.r.t. \mathcal{O} . Moreover, if \mathcal{D} is satisfiable w.r.t. \mathcal{O} and $q(\bar{x})$ is a Boolean or unary CQ of treewidth (ℓ', k') and $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$, then for all copies \bar{a}' of \bar{a} in $\mathcal{D}_{S, \ell, k}^{\approx}$:*

1. if $\mathcal{D}_{S, \ell, k}^{\approx}, \mathcal{O} \models q(\bar{a}')$, then $\mathcal{D}, \mathcal{O} \models q(\bar{a})$;
2. if $\mathcal{D}, \mathcal{O} \models q(\bar{a})$, $\ell' \leq \ell$, and $k' \leq k$, then $\mathcal{D}_{S, \ell, k}^{\approx}, \mathcal{O} \models q(\bar{a}')$;
3. if $\mathcal{D}, \mathcal{O} \models q(\bar{a})$ is witnessed by a homomorphism h from $q(\bar{x})$ to $\text{ch}_{\mathcal{O}}(\mathcal{D})$ that maps every variable distinct from \bar{x} to $\text{adom}(\text{ch}_{\mathcal{O}}(\mathcal{D})) \setminus \text{adom}(\mathcal{D})$, then $\mathcal{D}_{S, \ell, k}^{\approx}, \mathcal{O} \models q(\bar{a}')$.

Proof. We first consider Points 1., 2., and 3., and then return to satisfiability. Point 1 follows from the observation made above that there is a homomorphism from $\mathcal{D}_{S, \ell, k}^{\approx}$ to \mathcal{D} that maps \bar{a}' to \bar{a} . For Points 2. and 3., we first observe the following. Let $\vartheta = \phi(\bar{x}, \bar{y}) \rightarrow \exists \bar{z} \psi(\bar{x}, \bar{z})$ be any ℓ, k, ℓ', k' -TGD in \mathcal{O} and let $b \in \text{adom}(\mathcal{D})^{|\bar{x}|}$. Then, by Lemma 7 (2), there is a homomorphism from the body $\phi(\bar{x}, \bar{y})$ of ϑ to \mathcal{D} that maps \bar{x} to \bar{b} iff there is a homomorphism from $\phi(\bar{x}, \bar{y})$ to $\mathcal{D}_{S, \ell, k}^{\approx}$ that maps \bar{x} to \bar{b}' for every (equivalently: any) copy

\bar{b}' of \bar{b} . In terms of the construction of chase sequences it follows that ϑ is applicable to \mathcal{D} via a homomorphism that maps \bar{x} to \bar{b} iff it is applicable to $\mathcal{D}_{S, \ell, k}^{\approx}$ via a homomorphism that maps \bar{x} to \bar{b}' for every (equivalently: any) copy \bar{b}' of \bar{b} . It can thus be proved by induction over the construction of chase sequences that the following holds for the universal models $\text{ch}_{\mathcal{O}}(\mathcal{D})$ and $\text{ch}_{\mathcal{O}}(\mathcal{D}_{S, \ell, k}^{\approx})$: for every $b \in \text{adom}(\mathcal{D})$ and any copy b' of b in $\mathcal{D}_{S, \ell, k}^{\approx}$, the interpretation \mathcal{I}_b attached in $\text{ch}_{\mathcal{O}}(\mathcal{D})$ to \mathcal{D} at b is isomorphic to the interpretation $\mathcal{I}_{b'}$ attached in $\text{ch}_{\mathcal{O}}(\mathcal{D}_{S, \ell, k}^{\approx})$ to $\mathcal{D}_{S, \ell, k}^{\approx}$ at b' . Observe that the interpretations $\mathcal{I}_b, \mathcal{I}_{b'}$ have treewidth (ℓ', k') . Claim 3 follows directly from this observation. For Claim 2, use Lemma 7 (2).

It remains to show that \mathcal{D} is satisfiable w.r.t. \mathcal{O} iff $\mathcal{D}_{S, \ell, k}^{\approx}$ is satisfiable w.r.t. \mathcal{O} . The direction from left to right follows from the existence of a homomorphism from $\mathcal{D}_{S, \ell, k}^{\approx}$ to \mathcal{D} . Conversely, assume that $\mathcal{D}_{S, \ell, k}^{\approx}$ is satisfiable w.r.t. \mathcal{O} . Obtain an ontology \mathcal{O}' from \mathcal{O} by replacing all occurrences of \perp in \mathcal{O} by $\exists x A(x)$, for a single fresh concept name A . Then \mathcal{D} is satisfiable w.r.t. \mathcal{O}' (as any database is satisfiable w.r.t. \mathcal{O}') and \mathcal{D} is satisfiable w.r.t. \mathcal{O} iff $\mathcal{D}, \mathcal{O}' \not\models \exists x A(x)$. By Point 2 above, the latter is equivalent to $\mathcal{D}_{S, \ell, k}^{\approx}, \mathcal{O}' \not\models \exists x A(x)$ which holds since $\mathcal{D}_{S, \ell, k}^{\approx}$ is satisfiable w.r.t. \mathcal{O} . \square

Lemma 8 also justifies the notation $\langle \mathcal{D}_{\ell, k}^{\approx}, a \rangle$ denoting “the” database obtained from $\mathcal{D}_{\ell, k}^{\approx}$ by choosing a copy of a in $\mathcal{D}_{\ell, k}^{\approx}$ and renaming it back to a . In fact, one can now easily show that $\mathcal{D}_{\ell, k}^{\approx}, \mathcal{O} \models q(a')$ iff $\mathcal{D}_{\ell, k}^{\approx}, \mathcal{O} \models q(a'')$ for any possibly infinite set \mathcal{O} of TGDs, CQ $q(x)$, and any two copies a' and a'' of a constant a in \mathcal{D} .

Lemma 9. *Let $Q(\bar{x}) = (\mathcal{O}, \Sigma, q)$ be an OMQ with \mathcal{O} a possibly infinite ℓ, k, ℓ', k' -ontology, $1 \leq \ell < k$ and $1 \leq \ell' < k'$, and q a CQ of treewidth (ℓ, k) . Then for all Σ -databases \mathcal{D} and all $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$, $\bar{a} \in Q(\mathcal{D})$ iff $\bar{a} \in Q(\mathcal{D}_{\bar{a}, \ell, k}^{\approx})$.*

Proof. By Lemma 8, \mathcal{D} is satisfiable w.r.t. \mathcal{O} iff $\mathcal{D}_{\bar{a}, \ell, k}^{\approx}$ is satisfiable w.r.t. \mathcal{O} . Thus, we may assume that both \mathcal{D} and $\mathcal{D}_{\bar{a}, \ell, k}^{\approx}$ are satisfiable w.r.t. \mathcal{O} . If $\bar{a} \in Q(\mathcal{D}_{\bar{a}, \ell, k}^{\approx})$, then $\bar{a} \in Q(\mathcal{D})$ follows from the existence of a homomorphism from $\mathcal{D}_{\bar{a}, \ell, k}^{\approx}$ to \mathcal{D} mapping \bar{a} to \bar{a} . For the converse direction, let $\bar{a} \in Q(\mathcal{D})$. Let h be a homomorphism from q to $\text{ch}_{\mathcal{O}}(\mathcal{D})$ mapping \bar{x} to \bar{a} . Consider the restriction of h to the set of variables V in $q(x)$ that are mapped into $\text{adom}(\mathcal{D})$. Then, by Lemma 7 (1), we find a homomorphism g from V to the restriction of $\text{ch}_{\mathcal{O}}(\mathcal{D}_{\bar{a}, \ell, k}^{\approx})$ to $\text{adom}(\mathcal{D}_{\bar{a}, \ell, k}^{\approx})$ mapping \bar{x} to \bar{a} and such that $g(y)$ is a copy of $h(y)$, for all variables $y \in V$ (the latter condition is a consequence of the proof of Lemma 7 (1)). But then the remaining variables of q that are not in V can be mapped to the interpretations $\mathcal{I}_{g(y)}$ attached to $g(y)$ in $\text{ch}_{\mathcal{O}}(\mathcal{D}_{\bar{a}, \ell, k}^{\approx})$ in the same way they can be mapped to the interpretations $\mathcal{I}_{h(y)}$ attached to $h(y)$ in $\text{ch}_{\mathcal{O}}(\mathcal{D})$. The resulting homomorphism shows that $\bar{a} \in Q(\mathcal{D}_{\bar{a}, \ell, k}^{\approx})$. \square

B Proofs for Section 3

Lemma 2. Let $Q = (\mathcal{O}, \Sigma, q) \in (FO, UCQ)$. Then for all Σ -databases \mathcal{D} and $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$, $\bar{a} \in \text{app}_{\mathfrak{D}_\lambda}(Q, \mathcal{D})$ iff $\bar{a} \in Q(\mathcal{D}_{\bar{a}}^{\approx})$.

Proof. For the “if” direction, observe that for every finite subset \mathcal{D}' of $\mathcal{D}_{\bar{a}}^{\approx}$ containing \bar{a} the pointed database (\mathcal{D}', \bar{a}) is a member of \mathfrak{D}_λ . Moreover, there is a homomorphism from \mathcal{D}' to \mathcal{D} mapping \bar{a} to \bar{a} . Thus, this direction is a consequence of compactness and the preservation of OMQ answers under homomorphic images of databases (Bienvenu et al. 2014). For the “only if” direction, assume that there is a pointed tree database (\mathcal{D}', \bar{b}) with $\bar{b} \in Q(\mathcal{D}')$, and a homomorphism h from \mathcal{D}' to \mathcal{D} such that $h(\bar{b}) = \bar{a}$. To show that $\bar{a} \in Q(\mathcal{D}_{\bar{a}}^{\approx})$, it suffices to construct a homomorphism g from \mathcal{D}' to $\mathcal{D}_{\bar{a}}^{\approx}$ with $g(\bar{b}) = \bar{a}$ and to again use the preservation of OMQ answers under homomorphic images of databases. To define g , start with setting $g(\bar{b}) := h(\bar{b}) = \bar{a}$ and choose for any maximally connected component of the restriction of \mathcal{D}' to $\text{adom}(\mathcal{D}') \setminus \bar{b}$ an arbitrary constant c in that component and set $g(c) = h(c)$. Next, assume that $r(c_1, c_2) \in \mathcal{D}'$ with $c_1, c_2 \notin \bar{b}$, r a potentially inverse role, and $g(c_1)$ is already defined, but $g(c_2)$ is not. Then $g(c_1)rh(c_2)$ is a path in \mathcal{D} and we set $g(c_2) = g(c_1)rh(c_2)$. The restriction of g to $\text{adom}(\mathcal{D}') \setminus \bar{b}$ is a homomorphism by definition of $\mathcal{D}_{\bar{a}}^{\approx}$ and since, after removal of \bar{b} , \mathcal{D}' does not contain reflexive loops or multi-edges. To show that g is a homomorphism on \mathcal{D}' assume that $r(c_1, c_2) \in \mathcal{D}'$ with $c_1 \in \bar{b}$ and $c_2 \notin \bar{b}$, r a potentially inverse role. Then $\text{tail}(g(c_2)) = h(c_2)$ and so $(g(c_1), g(c_2)) \in \mathcal{D}_{\bar{a}}^{\approx}$ since $r(h(c_1), h(c_2)) \in \mathcal{D}$, as required. \square

C Proofs for Section 4

We begin by providing proofs that the examples given in this section are correct.

Example 3. Recall that in this example, $Q(x) = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALCC}, \text{CQ})$ with $\mathcal{O} = \{\top \sqsubseteq \forall r.(B_1 \rightarrow B) \sqcup \forall r.(B_2 \rightarrow B)\}$, $\Sigma = \{r, A, B, B_1, B_2\}$, and $q = \exists y r(x, y) \wedge A(y) \wedge B(y)$. Also,

$$\mathcal{D} = \{r(a, b_1), r(a, b_2), B_1(b_1), B_2(b_2), A(b_1), A(b_2)\}.$$

Let $Q_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u} = (\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u}, \Sigma, q)$ for $\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u}$ the result of removing from $\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u}$ all CIs that use a symbol that does not occur in \mathcal{O} . We claim that $a \notin Q_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u}(\mathcal{D})$ where $Q_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u} = (\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u}, \Sigma, q)$ for $\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u}$ the result of removing from $\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u}$ all CIs that use a symbol that does not occur in \mathcal{O} . To see this consider the universal model $\text{ch}_{\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u}}(\mathcal{D})$. Then $A(c) \in \text{ch}_{\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u}}(\mathcal{D})$ iff $c \in \{b_1, b_2\}$, but neither $B(b_1) \in \text{ch}_{\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u}}(\mathcal{D})$ nor $B(b_2) \in \text{ch}_{\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u}}(\mathcal{D})$ from which the claim follows. Observe that the latter holds also in $\text{ch}_{\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u}}(\mathcal{D})$, but in $\text{ch}_{\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u}}(\mathcal{D})$ there is an additional r -successor c of a with $A(c), B(c) \in \text{ch}_{\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u}}(\mathcal{D})$.

Example 4. Recall that in this example $Q = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALCC}, \text{CQ})$ is the Boolean OMQ with $\mathcal{O} = \{A \sqsubseteq B \sqcup$

$\forall r.B\}$, $\Sigma = \{r, A, B\}$, and $q = \exists x B(x)$. Assume $\mathcal{D} = \{A(a), r(a, b)\}$. We sketch the proof that $\mathcal{D} \not\models Q_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u}$. First observe that there does not exist any sequence of roles $\rho = r_1, \dots, r_n$ with $r_i \in \{r, r^-\}$ for all $1 \leq i \leq n$ such that $a \in Q'(\mathcal{D})$ for the query $Q' = (\mathcal{O}, \Sigma, q_\rho)$, where

$$q_\rho(x) = \exists x_1 \dots \exists x_n r_1(x, x_1) \wedge \dots \wedge r_n(x_{n-1}, x_n) \wedge B(x_n).$$

Thus, we find models \mathcal{I}_ρ of \mathcal{D} and \mathcal{O} such that $\mathcal{I}_\rho \not\models q_\rho(a)$. Consider the product $\mathcal{I} = \prod_\rho \mathcal{I}_\rho$. Then \mathcal{I} is a model of $\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u}$, as $\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u$ -ontologies are preserved under products (Hernich et al. 2020), and $\mathcal{I} \not\models q_\rho(a)$ for any ρ (we identify a with the sequence (a, a, \dots) in \mathcal{I}). Consider the subinterpretation \mathcal{I}' of \mathcal{I} induced by all nodes in \mathcal{I} that can be reached from a by any path ρ , that is, any path using r^+ and $(r^-)^+$. Then \mathcal{I}' is a model of \mathcal{D} and \mathcal{O} as \mathcal{O} does not use the universal role and $\mathcal{I}' \not\models \exists x B(x)$. It follows that $\mathcal{D} \not\models Q_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u}$, as required.

The following lemma is related to the notion of unraveling tolerance, introduced in (Lutz and Wolter 2017). In the language of that paper, Point 1 states that the DL $\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u$ is unraveling tolerant, even with infinite ontologies.

Note that for every database \mathcal{D} , $\text{adom}(\mathcal{D}) \subseteq \mathcal{D}_\emptyset^{\approx}$ by definition of tree unravelings. In contrast to unravelings into bounded treewidth, we thus need no ‘renaming back’ of those constants to their original name, c.f. the $\langle \mathcal{D}_{\ell, k}^{\approx}, \bar{a} \rangle$ notation from Section 6.

Lemma 10. Let $Q(\bar{x}) = (\mathcal{O}, \Sigma, q) \in (\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u, \text{bELIQ})$, \mathcal{D} a Σ -database, and $a \in \text{adom}(\mathcal{D})^{|\bar{x}|}$. Then

1. $\bar{a} \in Q(\mathcal{D})$ iff $\bar{a} \in Q(\mathcal{D}_\emptyset^{\approx})$;
2. \mathcal{D} is satisfiable w.r.t. \mathcal{O} iff $\mathcal{D}_\emptyset^{\approx}$ is satisfiable w.r.t. \mathcal{O} .

Proof. Point 1 is proved in (Lutz and Wolter 2017) for OMQs based on finite ontologies. The result easily lifts to infinite ontologies through compactness. Point 2 is a consequence of Point 1. To see the latter, let \mathcal{O}' be obtained from \mathcal{O} by replacing every CI $C \sqsubseteq \perp$ with $C \sqsubseteq A_\perp$, A_\perp a fresh concept name. Then \mathcal{D} is satisfiable w.r.t. \mathcal{O} iff $\mathcal{D}, \mathcal{O}' \models \exists u.A_\perp$ iff $\mathcal{D}_\emptyset^{\approx}, \mathcal{O}' \models \exists u.A_\perp$ (by Point 1) iff $\mathcal{D}_\emptyset^{\approx}$ is satisfiable w.r.t. \mathcal{O} . \square

The next lemma identifies a crucial property that relates $\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u$ -ontology relaxing approximation to tree unravelings of databases.

Lemma 3. Let $Q(\bar{x}) = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALCC}, \text{bELIQ})$, \mathcal{D} be a Σ -database, and $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$. Then

1. $\bar{a} \in Q_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u}(\mathcal{D})$ iff $\bar{a} \in Q(\mathcal{D}_\emptyset^{\approx})$;
2. \mathcal{D} is satisfiable w.r.t. $\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u}$ iff $\mathcal{D}_\emptyset^{\approx}$ is satisfiable w.r.t. \mathcal{O} .

Proof. We start with Point 1. First assume that $\bar{a} \in Q_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u}(\mathcal{D})$. By Lemma 10, $\bar{a} \in Q_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u}(\mathcal{D}_\emptyset^{\approx})$. From $\mathcal{O} \models \mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u}$, it follows that $\bar{a} \in Q(\mathcal{D}_\emptyset^{\approx})$.

Now assume that $\bar{a} \in Q(\mathcal{D}_\emptyset^{\approx})$. Then compactness yields a finite subset \mathcal{F} of $\mathcal{D}_\emptyset^{\approx}$ with $\bar{a} \in Q(\mathcal{F})$. We can view \mathcal{F} as an $\mathcal{E}\mathcal{L}\mathcal{I}$ -concept F . Then $\mathcal{O} \models F \sqsubseteq A$ and thus $F \sqsubseteq A$ is a CI in $\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u}$. If q is an ELIQ and $\bar{a} = a_0$, then we further have $\mathcal{D}_\emptyset^{\approx} \models F(a_0)$, thus $\mathcal{D}, \mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}_\perp^u} \models F(a_0)$ by

Lemma 10 which yields $\mathcal{D}, \mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}^u} \models A(a_0)$ as required. If q is a BELIQ $\exists u.C$, then there is an $a_0 \in \text{ind}(\mathcal{F})$ with $\mathcal{F}, \mathcal{O} \models C(a_0)$. Let F' be \mathcal{F} viewed as a BELIQ with root a_0 . Then $\mathcal{O} \models F' \sqsubseteq C$, thus $F' \sqsubseteq C$ is a CI in $\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}^u}$. From $\mathcal{D}_\emptyset^\approx \models F'(a_0)$, we obtain $\mathcal{D}, \mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}^u} \models F'(a_0)$ by Lemma 10 which yields $\mathcal{D}, \mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}^u} \models \exists u.C$ as required.

Now for Point 2. First assume that \mathcal{D} is unsatisfiable w.r.t. $\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}^u}$. By Lemma 10, $\mathcal{D}_\emptyset^\approx$ is unsatisfiable w.r.t. $\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}^u}$. From $\mathcal{O} \models \mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}^u}$, it thus follows that $\mathcal{D}_\emptyset^\approx$ is unsatisfiable w.r.t. \mathcal{O} .

Conversely, assume that \mathcal{D} is satisfiable w.r.t. $\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}^u}$. By Lemma 10, $\mathcal{D}_\emptyset^\approx$ is satisfiable w.r.t. $\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}^u}$. Further assume to the contrary of what is to be shown that $\mathcal{D}_\emptyset^\approx$ is unsatisfiable w.r.t. \mathcal{O} . Compactness yields a finite subset \mathcal{F} of $\mathcal{D}_\emptyset^\approx$ that is unsatisfiable w.r.t. \mathcal{O} . We can view \mathcal{F} as an $\mathcal{E}\mathcal{L}\mathcal{I}$ -concept F . Then $\mathcal{O} \models F \sqsubseteq \perp$ and thus $F \sqsubseteq \perp$ is a CI in $\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}^u}$, in contradiction of $\mathcal{D}_\emptyset^\approx$ being satisfiable w.r.t. $\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}^u}$. \square

We next prove Point 3 of Theorem 1. To prepare for the proof of Theorem 2 in the subsequent section, we actually establish a stronger result. By Point 1 of Lemma 3, we can decide whether $\bar{a} \in Q(\mathcal{D}_\emptyset^\approx)$ in place of $\bar{a} \in \mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{I}^u}(\mathcal{D})$. Here, consider the case $\bar{a} \in Q(\mathcal{D}_S^\approx)$ with $S \subseteq \text{adom}(\mathcal{D})$ such that S is given as an additional input with $|S| \leq |\bar{x}|$; we instantiate S with \emptyset for the proof of Theorem 1 and with \bar{a} for the proof of Theorem 2. Also, we consider ELIQs and disjunctions of BELIQs, a class of UCQs that we denote with BELIQ^\vee .

Theorem 9. *Given $Q(\bar{x}) = (\mathcal{O}, \Sigma, q) \in (\mathcal{AL}\mathcal{C}\mathcal{I}, \text{BELIQ}^\vee)$, a Σ -database \mathcal{D} , $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$, and $S \subseteq \text{adom}(\mathcal{D})$, it is decidable in time $2^{O(\|\mathcal{Q}\|)} \cdot O(\|\mathcal{D}\|)$ whether $\bar{a} \in Q(\mathcal{D}_S^\approx)$.*

We prove Theorem 9 by linear time reduction to the unsatisfiability of propositional Horn formulas, which is in linear time (Dowling and Gallier 1984). Let $Q(\bar{x}) = (\mathcal{O}, \Sigma, p)$, \mathcal{D} , \bar{a} , and S be as in the theorem.

It is convenient to view q as a concept, namely as an $\mathcal{E}\mathcal{L}\mathcal{I}$ -concept if q is an ELIQ and as a disjunction of $\mathcal{E}\mathcal{L}\mathcal{I}^u$ -concepts if q is a disjunction of BELIQs. We use $\text{c1}(\mathcal{O}, q)$ to denote the smallest set that contains all concepts in \mathcal{O} and the concept q and is closed under subconcepts and single negation (thus $\text{c1}(\mathcal{O}, q)$ contains $\mathcal{AL}\mathcal{C}\mathcal{I}$ -concepts and possibly negated $\mathcal{E}\mathcal{L}\mathcal{I}^u$ -concepts). A *type* for \mathcal{O} and q is a maximal set $t \subseteq \text{c1}(\mathcal{O}, q)$ such that $\prod t$ is satisfiable w.r.t. \mathcal{O} . We remark that this condition can be checked in time $2^{O(\|\mathcal{O}\|)}$. In fact, standard algorithms for the satisfiability of $\mathcal{AL}\mathcal{C}\mathcal{I}$ -concepts w.r.t. $\mathcal{AL}\mathcal{C}\mathcal{I}$ -ontologies such as type elimination (Baader et al. 2017) exhibit this running time. We use $\text{TP}(\mathcal{O}, q)$ to denote the set of all types for \mathcal{O} and q . The type *realized* at $d \in \Delta^{\mathcal{I}}$ in a model \mathcal{I} of \mathcal{O} is

$$\text{tp}_{\mathcal{I}}(d) := \{C \in \text{TP}(\mathcal{O}, q) \mid a \in C^{\mathcal{I}}\}.$$

For $t, t' \in \text{TP}(\mathcal{O}, q)$ and roles r , we write $t \rightsquigarrow_r t'$ if

- (c1) $D \in t'$ and $\exists r.D \in \text{c1}(\mathcal{O}, q)$ implies $\exists r.D \in t$ and
- (c2) $D \in t$ and $\exists r^-.D \in \text{c1}(\mathcal{O}, q)$ implies $\exists r^-.D \in t'$.

For a set $S' \subseteq \text{adom}(\mathcal{D})$, an *S' -assignment* is a function $\mu : S' \rightarrow \text{TP}(\mathcal{O}, q)$ such that the following conditions are satisfied for all $a, a_1, a_2 \in S'$:

- (c3) $A(a) \in \mathcal{D}$ implies $A \in \mu(a)$;
- (c4) $r(a_1, a_2) \in \mathcal{D}$ implies $\mu(a_1) \rightsquigarrow_r \mu(a_2)$;
- (c5) if q is an ELIQ and $\bar{a} = a$, then $q \notin \mu(a)$;
- (c6) if q is a disjunction of BELIQs, then $q \notin \mu(a)$.

We shall primarily be interested in S -assignments and in $S \cup \{a\}$ -assignments for constants $a \in \text{adom}(\mathcal{D}) \setminus S$. Note that there are at most $2^{O(\|\mathcal{Q}\|)}$ S -assignments and $S \cup \{a\}$ -assignments for each $a \in \text{adom}(\mathcal{D})$ and that we can compute the set of S -assignments and $S \cup \{a\}$ -assignments, $a \in \text{adom}(\mathcal{D})$, in time $2^{O(\|\mathcal{Q}\|)} \cdot \|\mathcal{D}\|$.

We introduce a propositional variable $p_{\mu,a}$ for every $a \in \text{adom}(\mathcal{D}) \setminus S$ and every $S \cup \{a\}$ -assignment μ , as well as a propositional variable p_μ for every S -assignment μ . Informally, $p_{\mu,a}$ being true means that it is not possible to simultaneously realize the type $\mu(a)$ at a (equivalently: at a path p with $\text{tail}(p) = a$) and the type $\mu(b)$ at b for every $b \in S$ in a model of \mathcal{D}_S^\approx and \mathcal{O} , and likewise for propositional variables p_μ . Now the propositional Horn formula φ is the conjunction of the following:

1. $\bigwedge_{\mu \in W} p_{\mu,a} \rightarrow p_{\mu',b}$ for all $r(a,b) \in \mathcal{D}$ with r a potentially inverse role and $a, b \notin S$ and all $S \cup \{b\}$ -assignments μ' , where W consists of all $S \cup \{a\}$ -assignments μ such that $\mu(a) \rightsquigarrow_r \mu'(b)$ and $\mu(c) = \mu'(c)$ for all $c \in S$;
2. $\bigwedge_{\mu \in W} p_{\mu,a} \rightarrow p_{\mu'}$ for all $a \in \text{adom}(\mathcal{D}) \setminus S$ and S -assignments μ' , where W consists of all $S \cup \{a\}$ -assignments μ such that $\mu(c) = \mu'(c)$ for all $c \in S$;
3. $p_\mu \rightarrow p_{\mu',a}$ for all $a \in \text{adom}(\mathcal{D}) \setminus S$, S -assignments μ , and $S \cup \{a\}$ -assignments μ' such that $\mu(c) = \mu'(c)$ for all $c \in S$;
4. $\bigwedge_{\mu \in W} p_{\mu,a} \rightarrow \perp$ for all $a \in \text{adom}(\mathcal{D}) \setminus S$ where W is the set of all $S \cup \{a\}$ -assignments.

It is clear that $\|\varphi\| \in 2^{O(\|\mathcal{Q}\|)} \cdot \|\mathcal{D}\|$ and that φ can be constructed in $O(2^{O(\|\mathcal{Q}\|)} \cdot \|\mathcal{D}\|)$. Now, Theorem 9 is an immediate consequence of the following.

Lemma 11. $\bar{a} \in Q(\mathcal{D}_S^\approx)$ iff φ is unsatisfiable.

Proof. The (contrapositive) of the “if” direction is proved as follows. Assume that $\bar{a} \notin Q(\mathcal{D}_S^\approx)$. Then there is a model \mathcal{I} of \mathcal{D}_S^\approx and \mathcal{O} with $\bar{a} \notin q(\mathcal{I})$. For each $a \in \text{adom}(\mathcal{D}) \setminus S$, let μ_a denote the $S \cup \{a\}$ -assignment that sets $\mu_a(b) = \text{tp}_{\mathcal{I}}(b)$ for all $b \in S \cup \{a\}$. Let μ_S be the S -assignment defined in the same way. It is readily checked that the μ_a and μ_S are indeed assignments and in particular, conditions (c5) and (c6) are satisfied since $\bar{a} \notin q(\mathcal{I})$.

Define a valuation V for the variables in φ by setting,

- $V(p_{\mu,a}) = 0$ iff $\mu = \mu_a$, for all $a \in \text{adom}(\mathcal{D}) \setminus S$;
- $V(p_\mu) = 0$ iff $\mu = \mu_S$.

It is readily checked that V is a model of φ .

For the (contrapositive of the) “if” direction, assume that φ is satisfied by some valuation V . Choose some

$a_0 \in \text{adom}(\mathcal{D}) \setminus S$ and some μ_0 with $V(p_{\mu_0, a_0}) = 0$, which exist due to the conjunct of φ in Point 4. For every $a \in \text{adom}(\mathcal{D}) \setminus S$, let T_a denote the set of types from $\text{TP}(\mathcal{O}, q)$ such that there is an $S \cup \{a\}$ -assignment μ such that $V(p_{\mu, a}) = 0$ and $\mu(b) = \mu_0(b)$ for all $b \in S$. Due to the conjuncts in Points 2 and 3, T_a is non-empty for all $a \in \text{adom}(\mathcal{D})$.

We construct a model \mathcal{I} of \mathcal{D}_S^{\approx} and \mathcal{O} such that $\text{tp}_{\mathcal{I}}(a) \in T_a(a)$ for all $a \in \text{adom}(\mathcal{D})$. By definition of T_a and due to Conditions (c5) and (c6) of assignments, this implies $\bar{a} \notin q(\mathcal{I})$ and thus it follows that $\bar{a} \notin Q(\mathcal{D}_S^{\approx})$, as required. We first assign to each $p \in \text{adom}(\mathcal{D}_S^{\approx})$ a type $t_p \in T_{\text{tail}(p)}$. Start with

- setting $t_a = \mu_0(a)$ for all $a \in S$ and
- and choosing $t_a \in T_a$ arbitrarily for all $a \in \text{adom}(\mathcal{D}) \setminus S$.

Now let $r(p_1, p_2) \in \mathcal{D}_S^{\approx}$, r a possibly inverse role, and assume that t_{p_1} is already defined, but t_{p_2} is not. By definition of \mathcal{D}_S^{\approx} , $p_2 = p_1 r a$ for some (possibly inverse) role r and some $a \in \text{adom}(\mathcal{D})$. Consequently, $r(\text{tail}(p_1), \text{tail}(p_2)) \in \mathcal{D}$. We know that $t_{p_1} \in T_{\text{tail}(p_1)}$. By the conjunct in Point 1, there is thus a $t_{p_2} \in T_{\text{tail}(p_2)}$ such that $t_{p_1} \rightsquigarrow_r t_{p_2}$.

We now construct an interpretation \mathcal{I} such that $\text{tp}_{\mathcal{I}}(p) = t_p$ for all $p \in \text{adom}(\mathcal{D}_S^{\approx})$. Start with setting

$$\begin{aligned} \Delta^{\mathcal{I}} &= \text{adom}(\mathcal{D}_S^{\approx}) \\ A^{\mathcal{I}} &= \{p \in \text{adom}(\mathcal{D}_S^{\approx}) \mid A \in t_p\} \\ r^{\mathcal{I}} &= \{(p_1, p_2) \mid r(p_1, p_2) \in \mathcal{D}_S^{\approx}\} \end{aligned}$$

for all concept names A and role names r . We then extend \mathcal{I} by adding, for all $p \in \text{adom}(\mathcal{D}_S^{\approx})$, a tree model \mathcal{I}_p of \mathcal{O} that satisfies type t_p at its root, with disjoint domain and identifying its root with p . It can be proved by induction on the structure of concepts D that for all concepts $D \in \text{cl}(\mathcal{O}, q)$,

1. $D^{\mathcal{I}} \cap \text{adom}(\mathcal{D}_S^{\approx}) = \{p \in \text{adom}(\mathcal{D}_S^{\approx}) \mid D \in t_p\}$ and
2. $D^{\mathcal{I}} \cap \Delta^{\mathcal{I}_p} = D^{\mathcal{I}_p} \cap \Delta^{\mathcal{I}_p}$ for all $p \in \text{adom}(\mathcal{D}_S^{\approx})$.

From this, it easily follows that \mathcal{I} is a model of \mathcal{O} and \mathcal{D}_S^{\approx} . \square

We now prove correctness of the careful chase algorithm.

Lemma 12. $\bar{a} \in Q_{\mathcal{E}\mathcal{L}\mathcal{T}_\perp^u}^{\approx}(\mathcal{D})$ iff $\bar{a} \in q(\mathcal{D}')$.

Proof. We start with the following observation.

Claim. The restriction of $\text{ch}_{\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{T}_\perp^u}^{\approx}}(\mathcal{D})$ to $\text{adom}(\mathcal{D})$ is identical to the restriction of \mathcal{D}' to $\text{adom}(\mathcal{D})$.

For a proof, it suffices to show that for all $a, b \in \text{adom}(\mathcal{D})$, the following conditions are satisfied:

- $r(a, b) \in \text{ch}_{\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{T}_\perp^u}^{\approx}}(\mathcal{D})$ iff $r(a, b) \in \mathcal{D}'$.

Immediate by definition of the chase and construction of \mathcal{D}' . In fact, $r(a, b) \in \text{ch}_{\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{T}_\perp^u}^{\approx}}(\mathcal{D})$ iff $r(a, b) \in \mathcal{D}$ iff $r(a, b) \in \mathcal{D}'$.

- $A(a) \in \text{ch}_{\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{T}_\perp^u}^{\approx}}(\mathcal{D})$ iff $A(a) \in \mathcal{D}'$.

From $A(a) \in \text{ch}_{\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{T}_\perp^u}^{\approx}}(\mathcal{D})$, it follows that $\mathcal{D}, \mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{T}_\perp^u}^{\approx} \models A(a)$ by Lemma 6 and thus $A(a) \in \mathcal{D}'$ by construction of \mathcal{D}' .

Conversely, let $A(a) \in \mathcal{D}'$. If $A(a) \in \mathcal{D}$, then $A(a) \in \text{ch}_{\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{T}_\perp^u}^{\approx}}(\mathcal{D})$ and we are done. Otherwise, $\mathcal{D}, \mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{T}_\perp^u}^{\approx} \models A(a)$ by construction of \mathcal{D}' . Since \mathcal{D} is satisfiable w.r.t. $\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{T}_\perp^u}^{\approx}$, we obtain $A(a) \in \text{ch}_{\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{T}_\perp^u}^{\approx}}(\mathcal{D})$.

Now for the proof of the lemma. For the ‘only if’ direction, assume that $\bar{a} \in Q_{\mathcal{E}\mathcal{L}\mathcal{T}_\perp^u}^{\approx}(\mathcal{D})$. By Lemma 6 and since \mathcal{D} is satisfiable w.r.t. $\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{T}_\perp^u}^{\approx}$, there is a homomorphism h from q to $\text{ch}_{\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{T}_\perp^u}^{\approx}}(\mathcal{D})$ with $h(\bar{x}) = \bar{a}$. It suffices to identify a homomorphism g from q to \mathcal{D}' with $g(\bar{x}) = \bar{a}$.

By the claim, we can start the definition of g by setting $g(x) = h(x)$ whenever $h(x) \in \text{adom}(\mathcal{D})$. For completing the definition, let g' be obtained from g in the following way:

- quantify all variables and then contract by identifying variables x, y whenever $h(x) = h(y)$;
- restricting to those atoms that contain at least one variable x with $h(x) \notin \text{adom}(\mathcal{D})$;
- for all remaining variables x with $h(x) \in \text{adom}(\mathcal{D})$: view x as an answer variable (drop quantification, if present), and re-add any atoms $A(x)$ and $r(x, x)$ from a .

It is clear that g' is a collection of ELIQs and BELIQs from $\text{trees}(q)$. It suffices to extend g to all these bELIQs.

Let $p(\bar{x})$ be a bELIQ in q' . Then $\mathcal{D}, \mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{T}_\perp^u}^{\approx} \models p(h(\bar{x}))$ is witnessed by h . By construction of \mathcal{D}' , a copy of p has been added to \mathcal{D}' with its root glued to $h(\bar{x})$ in case that p is an ELIQ and thus we can extend g to all variables in $p(\bar{x})$. If variables x_1 and x_2 have been identified during the construction of p with the resulting variable being x_1 , then additionally set $g(x_2) = g(x_1)$.

It is readily verified that g is indeed a homomorphism from q to \mathcal{D}' with $g(\bar{x}) = \bar{a}$.

For the ‘if’ direction, assume that $\bar{a} \in q(\mathcal{D}')$. Then there is a homomorphism h from q to \mathcal{D}' with $h(\bar{x}) = \bar{a}$. It suffices to show that there is a homomorphism g from \mathcal{D}' to $\text{ch}_{\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{T}_\perp^u}^{\approx}}(\mathcal{D})$ with $g(\bar{a}) = \bar{a}$ because then the composition of h with g yields a homomorphism from q to $\text{ch}_{\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{T}_\perp^u}^{\approx}}(\mathcal{D})$ that witnesses $\bar{a} \in Q_{\mathcal{E}\mathcal{L}\mathcal{T}_\perp^u}^{\approx}(\mathcal{D})$ by Lemma 6, as required.

By the claim, we can start by setting $g(a) = h(a)$ for all $a \in \text{adom}(\mathcal{D})$. Now assume that a copy of a bELIQ $p(\bar{x}) \in \text{trees}(q)$ has been added during the construction of \mathcal{D}' , glueing the root to some $a \in \text{adom}(\mathcal{D})$ if p is an ELIQ. Let $\bar{a} = a$ if p is an ELIQ and $\bar{a} = ()$ otherwise. Then $\mathcal{D}, \mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{T}_\perp^u}^{\approx} \models p(\bar{a})$. By Lemma 6 and since \mathcal{D} is satisfiable w.r.t. $\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{T}_\perp^u}^{\approx}$, we find a homomorphism h_p from $p(\bar{x})$ to $\text{ch}_{\mathcal{O}_{\mathcal{E}\mathcal{L}\mathcal{T}_\perp^u}^{\approx}}(\mathcal{D})$ with $h_p(\bar{x}) = \bar{a}$. We can extend g to the added copy of $p(\bar{x})$ based on h_p . \square

It remains to analyze the running time of the careful chase algorithm. The computation of \mathcal{D}' needs time $2^{O(\|q\|^2)} \cdot |\text{adom}(\mathcal{D})| \cdot 2^{O(\|Q\|)} \cdot O(\|\mathcal{D}\|)$ while the size of \mathcal{D}' is bounded by $\|\mathcal{D}\| + |\text{adom}(\mathcal{D})| \cdot 2^{O(\|q\|^2)}$. The unsatisfiability check in (i) runs in time $2^{O(\|Q\|)} \cdot O(\|\mathcal{D}\|)$. Checking $\bar{a} \in q(\mathcal{D}')$ in (ii) by brute force means to test $\|\mathcal{D}'\|^{|\bar{q}|}$ candidate homomorphisms, that is, single exponentially many in the case of combined complexity and polynomially many in the case of data complexity, which gives the EXPTIME and PTIME upper bounds in Point 1. If $q \in \text{UCQ}_k^{\text{tw}}$, we can check $\bar{a} \in q(\mathcal{D}')$ using a blackbox procedure that runs in PTIME combined complexity (Dechter and Pearl 1989; Freuder 1990), which yields the FPT upper bound in Point 2.

D Proofs for Section 5

We prove the upper bounds in Theorem 2.

Theorem 2. *Tree-database relaxing OMQ evaluation is*

1. 2EXPTIME-complete in combined complexity and in linear time in data complexity (thus FPT) with running time $2^{2^{O(\|Q\|)}} \cdot O(\|\mathcal{D}\|)$ in $(\mathcal{ALCCI}, \mathcal{Q})$, $\mathcal{Q} \in \{\text{CQ}, \text{UCQ}, \text{CQ}_k^{\text{tw}}, \text{UCQ}_k^{\text{tw}} \mid k \geq 1\}$;
2. EXPTIME-complete in combined complexity and in linear time in data complexity (thus FPT) with running time $2^{O(\|Q\|)} \cdot O(\|\mathcal{D}\|)$ in $(\mathcal{ALC}, \mathcal{Q})$ and $(\mathcal{ALCCI}, \mathcal{Q})$, $\mathcal{Q} \in \{\text{AQ}, \text{BELIQ}\}$.

Let $Q(\bar{x}) = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALCCI}, \text{UCQ})$, \mathcal{D} be a Σ -database, and $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$. Let q^c be obtained from q by replacing each CQ p with the UCQ that consists of all contractions p' of p such that the restriction of p' to the quantified variables in it is a tree (recall that, for us, trees need not be connected). The number of CQs in q^c is bounded by $2^{\|q\|^2}$ and the size of each CQ is bounded by $\|q\|$. Let $Q^c = (\mathcal{O}, \Sigma, q^c)$.

Lemma 13. $\bar{a} \in Q(\mathcal{D}_{\bar{a}}^{\approx})$ iff $\bar{a} \in Q^c(\mathcal{D}_{\bar{a}}^{\approx})$.

Proof. The “if” direction is immediate. For “only if”, assume that $\bar{a} \notin Q^c(\mathcal{D}_{\bar{a}}^{\approx})$. Then there is a forest model \mathcal{I} of $\mathcal{D}_{\bar{a}}^{\approx}$ and \mathcal{O} such that $\bar{a} \notin q^c(\mathcal{I})$, that is, $G_{\mathcal{I}}$ is obtained from $G_{\mathcal{D}_{\bar{a}}^{\approx}}$ by attaching (possibly infinite) trees at constants from $\text{adom}(\mathcal{D}_{\bar{a}}^{\approx})$ (Lutz 2008a). It suffices to show that $\bar{a} \notin q(\mathcal{I})$. Assume to the contrary that this is the case. Then there is a homomorphism h from a CQ p in q to \mathcal{I} with $h(\bar{x}) = \bar{a}$. Let p' be obtained from p by identifying variables x and y if $h(x) = h(y)$ and at least one of x, y is not an answer variable. Using the definition of $\mathcal{D}_{\bar{a}}^{\approx}$ and of q^c and the fact that \mathcal{I} is a forest model, it is straightforward to show that p' is a CQ in p^c . Moreover, h is a homomorphism from p' to \mathcal{I} , in contradiction to $\bar{a} \notin q^c(\mathcal{I})$. \square

We can thus work with Q^c in place of Q .

We next decorate \mathcal{D} and decorate as well as simplify q^c . Let $\bar{x} = x_0 \cdots x_{n-1}$ and $\bar{a} = a_0 \cdots a_{n-1}$. For each (possibly inverted) role name r used in q and all $i < n$, introduce a fresh concept name A_{r,a_i} that, informally, expresses the existence of an r -edge to a_i . Extend \mathcal{D} to a database \mathcal{D}^d by adding $A_{r,a_i}(b)$ for all $r(b, a_i) \in \mathcal{D}$. This is clearly possible in time $O(\|\mathcal{D}\| \cdot \|q\|^2)$. Further, let q^d be obtained from q^c by doing the following for each CQ p in q^c :

1. replace each atom $r(y, x_i)$, y a quantified variable, with $A_{r,a_i}(y)$;
2. for each atom $r(x_{i_1}, x_{i_2})$, $0 \leq i_1, i_2 < n$, check whether $r(a_{i_1}, a_{i_2}) \in \mathcal{D}$; if the check fails, remove p from the UCQ; if it succeeds, remove atom $r(x_{i_1}, x_{i_2})$ from p ;
3. for each atom $A(x_i)$, $0 \leq i < n$, check whether $\mathcal{D}_{\bar{a}}^{\approx}, \mathcal{O} \models A(a_i)$ using Theorem 9 with $S = \bar{a}$; if the check fails, remove p from the UCQ; if it succeeds, remove atom $A(x_i)$ from p .

If some of the CQs in q^c becomes empty in the process (that is, all of its atoms have been dropped), then we know that $\bar{a} \in Q^c(\mathcal{D}_{\bar{a}}^{\approx})$ and return ‘yes’. If all CQs have been dropped, then we know that $\bar{a} \notin Q^c(\mathcal{D}_{\bar{a}}^{\approx})$ and return ‘no’. Clearly all of the above can be done in time $2^{O(\|Q\|^2)} \cdot O(\|\mathcal{D}\|)$. Note that q^d contains no answer variables as all atoms that mention them have been dropped. It is not hard to show the following using some basic manipulations of homomorphisms that witness query answers.

Lemma 14. $\bar{a} \in Q^c(\mathcal{D}_{\bar{a}}^d)$ iff $(\mathcal{D}_{\bar{a}}^d)^{\approx}, \mathcal{O} \models q^d$.

It thus remains to decide whether $(\mathcal{D}_{\bar{a}}^d)^{\approx}, \mathcal{O} \models q^d$. To ease notation, from now on we write \mathcal{D} instead of \mathcal{D}^d . By construction, q^d is a UCQ in which each CQ is a disjoint union of BELIQs. In other words, it is a disjunction of conjunctions of BELIQs. Using the laws of distributivity, we can convert it into an equivalent conjunction of disjunctions of BELIQs $q_1 \wedge \cdots \wedge q_k$. To decide whether $\mathcal{D}_{\bar{a}}^{\approx}, \mathcal{O} \models q^d$, it suffices to decide whether $\mathcal{D}_{\bar{a}}^{\approx}, \mathcal{O} \models q_i$ for $1 \leq i \leq k$. This can be done in time $2^{O(\|Q\| + \|q^d\|)} \cdot O(\|\mathcal{D}\|)$ by invoking Theorem 9 with $S = \bar{a}$.

If $Q \in (\mathcal{ALCCI}, \text{UCQ})$, then the above procedure runs in time $2^{2^{O(\|Q\|^2)}} \cdot O(\|\mathcal{D}\|)$, which yields both the 2EXPTIME upper bound in combined complexity and the linear time upper bound in data complexity in Point 1 of Theorem 2. If $Q \in (\mathcal{ALCCI}, \text{BELIQ})$, then we can omit the initial construction of q^c and, as a consequence, obtain a running time of $2^{O(\|Q\|)} \cdot O(\|\mathcal{D}\|)$ and thus the EXPTIME upper bound in combined complexity as well as the linear time upper bound in data complexity in Point 2 of Theorem 2.

Theorem 3. *Tree-database relaxing OMQ evaluation is EXPSPACE-hard in $(\mathcal{ALC}, \text{CQ})$ and 2EXPTIME-hard in $(\mathcal{ALC}, \text{UCQ})$.*

Proof. In (Lutz 2008a), it is shown that evaluating OMQs from $(\mathcal{ALCCI}, \text{CQ})$ is 2EXPTIME-hard on databases of the form $\{A_0(a)\}$ and for Boolean OMQs that use only a single role name r . The proof is by reduction from the word problem for exponentially space bounded alternating Turing machines (ATMs). It is not difficult to modify the reduction so that it uses exponentially time bounded ATMs instead, thus only showing EXPSPACE-hardness.

We thus know that the following problem is EXPSPACE-hard: given a Boolean OMQ $Q = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALCCI}, \text{CQ})$ with a single role name r and a concept name $A_0 \in \Sigma$, decide whether there is a tree-shaped model \mathcal{I} of \mathcal{O} with A_0 true at the root and $\mathcal{I} \models q$. We call such an \mathcal{I} a *tree witness* for Q and A_0 . An inspection of the constructions in (Lutz

2008a) shows that we can concentrate on witness trees that have depth at most $2^n - 2$, n the size of Q , and in which every node has at most 5 neighbors.² We refer to such an \mathcal{I} as a *small tree witness*.

Let $Q = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALCC}, \text{CQ})$ be Boolean with single role name r and $A_0 \in \Sigma$. We construct in polynomial time a Boolean OMQ $Q' = (\mathcal{O}', \Sigma, q') \in (\mathcal{ALCC}, \text{CQ})$ and Σ -database \mathcal{D} such that there is a tree witness for Q and A_0 if and only if $\mathcal{D} \approx, \mathcal{O}' \not\models q$. This gives the desired EXPSPACE lower bound.

We start with the construction of the database \mathcal{D} . The idea is to design $\mathcal{D} \approx$ so that small tree witnesses \mathcal{I} for Q and A_0 can be embedded into models \mathcal{J} of $\mathcal{D} \approx$ whose domain is $\text{adom}(\mathcal{D} \approx)$, that is, when no new elements are introduced by existential quantification. The constants used in \mathcal{D} take the form

$$a_{R,i,j} \text{ with } R \in \{r, r^-\}, 0 \leq i \leq 4, \text{ and } 0 \leq j \leq 2.$$

Informally, R indicates that copies of $a_{R,i,j}$ in the unraveling serve as R -successors, i is used to achieve that every constant in $\mathcal{D} \approx$ has 5 r -successors and 5 r^- -successors (note that the 5 neighbors in small tree witnesses could be linked via r or r^- , hence we prepare for both), and j is used to implement a ‘directionality’ in unravelings. Now, \mathcal{D} contains the following facts, for all $R \in \{r, r^-\}$, $0 \leq i, i' \leq 4$, and $0 \leq j \leq 2$, and where $j \oplus 1$ denotes $j + 1$ modulo 5:

- $r(a_{R,i,j}, a_{r,i',j \oplus 1})$ and $r(a_{r^-,i',j \oplus 1}, a_{R,i,j})$;
- $A_{R,i,j}(a_{R,i,j})$.

Recall that the constants of unraveled databases are paths. In Figure 4, we show the neighborhood of a path in $\mathcal{D} \approx$ that ends with a constant of the form $a_{r,i,j}$. All edges are r -edges. The neighborhood of paths ending with constants of the form $a_{r^-,i',j}$ is similar, except that the direction of the edge from above and the right half of the edges from below is reversed.

We next assemble the ontology \mathcal{O}' . We would like to say that the roots of $\mathcal{D} \approx$ are labeled with A_0 (or at least one of them is) while other ‘copies’ of the same constant from \mathcal{D} are not. However, there seems no way to do this since all such copies are bisimilar in $\mathcal{D} \approx$. This issue is in a sense the only obstacle in improving our EXPSPACE lower bound to a 2EXPTIME one.

To address the issue, we install a binary counter that counts the depth of constants in $\mathcal{D} \approx$ modulo 2^n . Every counter value is reached infinitely often and we make A_0 true at every constant with counter value 0. This allows us to embed infinitely many copies of a small tree witness into $\mathcal{D} \approx$. Note that the small tree witness has depth $2^n - 2$, and thus the depth $2^n - 1$ is present as a counter value in $\mathcal{D} \approx$, but not needed for representing nodes of the witness. We are instead going to use the constants with counter value $2^n - 1$ to separate different copies of the tree witness from each other. This is done via a concept name X that identifies

²This is why we switch to exponentially time bounded ATMs; for exponentially space bounded ones, models can become double exponentially deep and it is not clear how to deal with that in the remainder of the reduction.

those constants that ‘exist’ in the sense that they correspond to a domain element of the small tree witness.

We introduce concept names L_0, \dots, L_{n-1} to implement the counter. For every $R \in \{r, r^-\}$, $i \in \{0, 1, 2, 4\}$, $j \in \{0, 1, 2\}$ and $k \in \{0, \dots, n-1\}$, \mathcal{O}' contains the following CIs:

$$\begin{aligned} A_{R,i,j} \sqcap \prod_{0 \leq i < k} L_i \sqcap \neg L_k &\sqsubseteq \forall r.(A_{r,i',j \oplus 1} \rightarrow L_k) \\ A_{R,i,j} \sqcap \prod_{0 \leq i \leq k} L_i &\sqsubseteq \forall r.(A_{r,i',j \oplus 1} \rightarrow \neg L_k) \\ A_{R,i,j} \sqcap \bigsqcup_{0 \leq i < k} \neg L_i \sqcap \neg L_k &\sqsubseteq \forall r.(A_{r,i',j \oplus 1} \rightarrow \neg L_k) \\ A_{R,i,j} \sqcap \bigsqcup_{0 \leq i < k} \neg L_i \sqcap L_k &\sqsubseteq \forall r.(A_{r,i',j \oplus 1} \rightarrow L_k) \\ \exists r.(A_{R,i,j} \sqcap \prod_{0 \leq i < k} L_i \sqcap \neg L_k) &\sqsubseteq A_{r^-,i',j \oplus 1} \rightarrow L_k \\ \exists r.(A_{R,i,j} \sqcap \prod_{0 \leq i \leq k} L_i) &\sqsubseteq A_{r^-,i',j \oplus 1} \rightarrow \neg L_k \\ \exists r.(A_{R,i,j} \sqcap \bigsqcup_{0 \leq i < k} \neg L_i \sqcap \neg L_k) &\sqsubseteq A_{r^-,i',j \oplus 1} \rightarrow \neg L_k \\ \exists r.(A_{R,i,j} \sqcap \bigsqcup_{0 \leq i < k} \neg L_i \sqcap L_k) &\sqsubseteq A_{r^-,i',j \oplus 1} \rightarrow L_k \end{aligned}$$

Note that the last four lines essentially serve the same purpose as the first four ones. We have to use a different formulation, though, as we do not want to use quantification over inverse roles to stay within \mathcal{ALCC} . We next say that depth 0 corresponds to the root of tree witnesses:

$$\neg L_0 \sqcap \dots \sqcap \neg L_{n-1} \sqsubseteq A_0 \sqcap X$$

The rest of the ontology \mathcal{O}' is an adaptation of \mathcal{O} . In particular, we replace existential quantification with universal quantification, exploiting the fact that unraveling has already generated objects that can serve as witnesses for existential quantifiers. This also allows us to stay within \mathcal{ALCC} despite the fact that the original ontology \mathcal{O} is formulated in \mathcal{ALCC} .

We can assume w.l.o.g. that \mathcal{O} has the form $\{\top \sqsubseteq C_{\mathcal{O}}\}$ with $C_{\mathcal{O}}$ in negation normal form, that is, negation is only applied to concept names, but not to compound concepts. Introducing fresh concept names, it is straightforward to further transform \mathcal{O} so that it consists only of CIs of the form $\top \sqsubseteq A$, $A \sqsubseteq B$, $A \sqsubseteq \neg B$, $A \sqsubseteq B_1 \sqcup B_2$, $A \sqsubseteq \exists R.B$, and $A \sqsubseteq \forall R.B$ where A, B, B_1 , and B_2 are all concept names and $R \in \{r, r^-\}$. The fresh concept names are of course not included in the database schema, so the resulting OMQ is equivalent to the original one. We reflect the CIs from \mathcal{O} in \mathcal{O}' as follows:

- CIs $\top \sqsubseteq A$, $A \sqsubseteq B$, $A \sqsubseteq \neg B$, $A \sqsubseteq B_1 \sqcup B_2$ remain as they are except that we conjunctively add X on the left-hand side;
- CIs $A \sqsubseteq \forall r.B$, become $X \sqcap A \sqsubseteq \forall r.(X \rightarrow B)$;
- CIs $A \sqsubseteq \forall r^-.B$, become $X \sqcap \exists r.(A \sqcap X) \sqsubseteq B$;
- CIs $A \sqsubseteq \exists R.B$, $R \in \{r, r^-\}$, are translated into the following CIs for all $i \in \{0, \dots, 4\}$ and $j \in \{0, 1, 2\}$:

$$\begin{aligned} X \sqcap A \sqcap A_{R,i,j} &\sqsubseteq \bigsqcup_{0 \leq i' \leq 4} \forall r.(A_{R,i',j \oplus 1} \rightarrow (X \sqcap B)) \\ X \sqcap A \sqcap A_{R^-,i',j} &\sqsubseteq \bigsqcup_{0 \leq i' \leq 4} \forall r.(A_{R,i',j \oplus 1} \rightarrow (X \sqcap B)). \end{aligned}$$

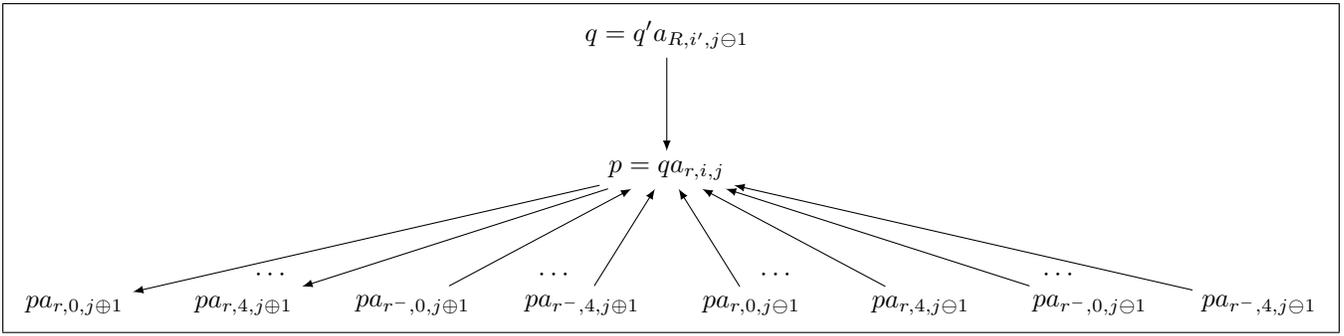


Figure 4: Constant $qa_{r,i,j}$ and its neighborhood in \mathcal{D}^\approx .

This finishes the construction of the ontology \mathcal{O}' . The CQ q' is obtained from q by adding the atom $X(x)$ for every variable x in q . Let $Q' = (\mathcal{O}', \Sigma, q')$. It remains to show the following.

Claim. There is a small tree witness for Q and A_0 if and only if $\mathcal{D}^\approx, \mathcal{O}' \not\models q'$.

Let \mathcal{I} be a small tree witness for Q and A_0 . We construct a model \mathcal{J} of \mathcal{D}^\approx and \mathcal{O}' that witnesses $\mathcal{D}^\approx, \mathcal{O}' \not\models q'$. Intuitively, \mathcal{J} is obtained by representing infinitely many copies of \mathcal{I} in \mathcal{D}^\approx , where elements of $\text{adom}(\mathcal{I})$ are represented by constants that satisfy X . To construct \mathcal{J} , we start with \mathcal{D}^\approx seen as an interpretation and then extend concept extensions, but we do neither extend role assertions nor add new constants. First, we assign counter values to all constants in $\Delta^\mathcal{J}$. This can be done by starting with the roots, that is, with all paths in $\Delta^\mathcal{J}$ that consist of a single constant. Assigning to all such $a \in \Delta^\mathcal{J}$ the counter value zero, thus $a \notin L_i^\mathcal{J}$ for $0 \leq i < n$. Then inductively assign incremented counter values (modulo 2^n) to all neighbours of constants that have already been assigned a counter value.

Next, we embed \mathcal{I} into \mathcal{J} infinitely often. Let $Z \subseteq \Delta^\mathcal{J}$ be the set of elements that have been assigned counter value zero. For every $a \in Z$, inductively construct an injective function h_a from $\Delta^\mathcal{I}$ to $\Delta^\mathcal{J}$ as follows:

- h_a maps the root of \mathcal{I} to a ;
- if $h_a(b) = c \in A_{R,i,j}^\mathcal{J}$ and $(b, b') \in r^\mathcal{I}$ with $h_a(b')$ undefined, then choose as $h_a(b')$ an r -successor of $h_a(b)$ in \mathcal{J} that is in $A_{r,i',j\oplus 1}^\mathcal{J}$ for some i' ;
- if $h_a(b) = c \in A_{R,i,j}^\mathcal{J}$ and $(b', b) \in r^\mathcal{I}$, then choose as $h_a(b')$ an r -successor of $h_a(b)$ in \mathcal{J} that is in $A_{r^-,i',j\oplus 1}^\mathcal{J}$ for some i' .

Since every node in \mathcal{I} has at most 5 neighbors, such an injective function exists; and since the depth of \mathcal{I} is at most $2^n - 1$, h_a never hits an element of \mathcal{J} whose counter value is $2^n - 1$. We extend \mathcal{J} with concept names such that every h_a becomes a homomorphism, and additionally set $X^\mathcal{J} = \bigcup_{a \in Z} \text{range}(h_a)$. This finishes the definition of \mathcal{J} . It can be checked that \mathcal{J} is indeed a model of \mathcal{O}' .

For the other direction, let $\mathcal{D}^\approx, \mathcal{O}' \not\models q'$, so there is a model \mathcal{J} of \mathcal{D}^\approx and \mathcal{O}' with $\mathcal{J} \not\models q'$. By construction of

\mathcal{O}' , we can assume w.l.o.g. that $\Delta^\mathcal{J} = \text{adom}(\mathcal{D}^\approx)$. Since \mathcal{O}' enforces a counter that counts modulo 2^n , there are infinitely many elements $a \in \Delta^\mathcal{J}$ such that $a \notin L_i^\mathcal{J}$ for $0 \leq i < n$. From now on, let a be a fixed such element. We construct a tree witness \mathcal{I} for Q and A_0 . Let $\Delta^\mathcal{I}$ be the smallest subset of $\Delta^\mathcal{J}$ that contains a and such that if $b \in \Delta^\mathcal{I}$, $(b, b') \in r^\mathcal{J} \cup (r^-)^\mathcal{J}$, and $b' \in X^\mathcal{J}$, then $b' \in \Delta^\mathcal{I}$. For every symbol $\alpha \in \text{sig}(\mathcal{O}) \cup \Sigma$, let $\alpha^\mathcal{I}$ be the restriction of $\alpha^\mathcal{J}$ to $\Delta^\mathcal{I}$. It can be checked that \mathcal{I} is indeed a tree witness for Q and A_0 .

We only sketch the idea of the 2EXPTIME lower bound for $(\mathcal{ALC}, \text{UCQ})$. We again start from the result established in (Lutz 2008a) that evaluating Boolean OMQs from $(\mathcal{ALC}, \text{CQ})$ is 2EXPTIME-hard on databases of the form $\{A_0(a)\}$. Unlike for the EXPSpace version, however, tree witnesses now have double exponential depth. This is a main obstacle to re-using the above proof because we cannot implement a double exponential counter using the ontology. The solution is to implement it using the query instead, which for this purpose we need to be a UCQ.

Let us get into some more detail. We implement a binary counter that counts modulo 2^{2^n} by attaching a binary subtree of depth n below each of the already existing nodes of \mathcal{D}^\approx . The 2^n leaves of the tree serve as the bit positions of the counter and concept names T and F are used to signal the value of each bit. We can make room for the additional trees by letting i range from 0 to 6 instead of from 0 to 4.

The ontology \mathcal{O}' is constructed to ensure that exactly one of T and F is true at each leaf of a counting tree. It then suffices to construct a Boolean CQ p such that for any model \mathcal{J} of \mathcal{O}' and \mathcal{D}^\approx with $\Delta^\mathcal{J} = \text{adom}(\mathcal{D}^\approx)$, $\mathcal{J} \models p$ iff somewhere in \mathcal{J} the counter is not properly incremented. For if we have found such a p , we can use $p' \vee q$ in place of q' in the reduction presented above.

Now, the (non-trivial) CQs constructed in the mentioned proof in (Lutz 2008a) are exactly what is needed for p . While there are no counting trees in (Lutz 2008a), there are configuration trees for exponentially space bounded ATMs, and identifying incrementation defects in counting trees is essentially the same task as identifying defects in ATM transitions (the former is actually slightly simpler). We refrain from repeating the details of the construction of p and only remark that p uses an alternation of inverse r -edges and non-inverse r -edges, and thus the edges of our counting trees

need to consist of an inverse r -edge followed by an r -edge (and thus these trees actually have depth $2n$ rather than depth n , but branch only on every second level). \square

E Proofs for Section 6

Proposition 2. $\omega, \omega, 1, 2$ -ontology relaxing OMQ evaluation in $(\mathcal{ALC}, \text{AQ})$ is CONP-hard in data complexity and there are ℓ', k' such that $1, 2, \ell', k'$ -ontology relaxing OMQ evaluation in $(\mathcal{ALCI}, \text{CQ})$ is 2EXPTIME-hard in combined complexity.

Proof. CONP-hardness is shown as in Example 7: consider the Boolean OMQ $Q(\mathcal{O}, \{e\}, \exists x D(x))$ given there that expresses non-3-colorability. Then $\mathcal{O}_{\omega, \omega, 1, 1}^{\approx}$ contains the TGD $q_{\mathcal{D}} \rightarrow \exists x D(x)$ for every undirected graph \mathcal{D} that is not 3-colorable. Thus, $\mathcal{O}_{\omega, \omega, 1, 1}^{\approx}$ is equivalent to Q .

For 2EXPTIME-hardness, we recall the result from (Lutz 2008a) that evaluating OMQs from $(\mathcal{ALCI}, \text{CQ})$ is 2EXPTIME-hard. It is easy to verify that the proof uses only Boolean OMQs, CQs of bounded treewidth, and databases of the simple form $\{A(a)\}$. In other words, there are ℓ', k' such that given an OMQ $Q = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALCI}, \text{CQ})$ with q of treewidth (ℓ', k') and a database of the form $\mathcal{D} = \{A(a)\}$, it is 2EXPTIME-hard to decide whether $\mathcal{D} \models Q$. This, however, is the case if and only if $\mathcal{D} \models Q_{\omega, \omega, 1, 1}^{\approx} = (\mathcal{O}_{\omega, \omega, 1, 1}^{\approx}, \Sigma, q)$ since $\mathcal{D} \models Q$ implies that $\mathcal{O}_{\omega, \omega, 1, 1}^{\approx}$ contains the TGD $A(x) \rightarrow q()$. \square

We now turn to Example 8, providing the missing proof that $\mathcal{D} \not\models Q_{\omega, \omega, 1, m^2-1}^{\approx}$.

It is sufficient to argue that $\mathcal{U} \not\models q_m$, where for brevity $\mathcal{U} = \text{ch}_{\mathcal{O}_{\omega, \omega, 1, m^2-1}^{\approx}}(\mathcal{D})$. Assume the contrary that there is a homomorphism h from q_m to \mathcal{U} . By construction, \mathcal{U} has treewidth $(1, m^2 - 1)$. The query q_m has treewidth $(1, m^2)$, but not treewidth $(1, m^2 - 1)$, so h is not injective, say $h(x_{i_0, j_0}) = h(x_{i_1, j_1})$ for some $i_0, j_0, i_1, j_1 \in \{1, \dots, m\}$ with $(i_0, j_0) \neq (i_1, j_1)$. Since h is a homomorphism, $h(x_{i_0, j_0}) \in A_{i_0, j_0}^{\mathcal{U}} \cap A_{i_1, j_1}^{\mathcal{U}}$. By analysing the TGDs in $\mathcal{O}_{\omega, \omega, 1, m^2-1}^{\approx}$, however, it is clear that \mathcal{U} does not contain any element that satisfies both A_{i_0, j_0} and A_{i_1, j_1} .

Proposition 3. Let $\ell, k \geq 1$ with $\ell < k$. For every $Q(\bar{x}) = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALCI}, \text{UCQ})$ and $k' \geq |\text{var}(q)|$, $\mathcal{Q}_{\ell, k, 1, |\text{var}(q)|}^{\approx} \equiv \mathcal{Q}_{\ell, k, 1, k'}^{\approx}$.

Proof. The proposition can be proved by analyzing the careful chase algorithm presented below. In fact, it is a consequence of Lemma 17 and the fact that $\text{trees}^{1, k'} = \text{trees}^{1, |\text{var}(q)|}$ whenever $k' \geq |\text{var}(q)|$ which implies that the constructed database \mathcal{D}' is exactly identical for all such k' . \square

We now turn to giving a detailed description of the careful chase algorithm and proving its correctness, with the aim of establishing the upper bounds in Theorem 4. For doing so, we assume that Theorem 5 is already available. A proof of the latter is provided in the subsequent section. We start with the observation that it suffices to treat CQs in place of

UCQs. The following is a consequence of the fact that TGD-ontologies have universal models, even if infinite.

Lemma 15. Let $Q(\bar{x}) = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALCI}, \text{UCQ})$, $q = p_1 \vee \dots \vee p_n$, \mathcal{D} a Σ -database, and $\ell, k, \ell', k' \geq 1$ with $\ell < k$ and $\ell' < k'$. Then $\mathcal{Q}_{\ell, k, \ell', k'}^{\approx}(\mathcal{D}) = Q_1(\mathcal{D}) \cup \dots \cup Q_n(\mathcal{D})$, $Q_i = (\mathcal{O}_{\ell, k, \ell', k'}^{\approx}, \Sigma, p_i)$ for $1 \leq i \leq n$.

Together with Theorem 5, the following lemma allow us to check satisfiability of a database \mathcal{D} w.r.t. the infinite ontology $\mathcal{O}_{\ell, k, \ell', k'}^{\approx}$, given only \mathcal{D} and \mathcal{O} .

Lemma 16. For any database \mathcal{D} and \mathcal{ALCI} -ontology \mathcal{O} , \mathcal{D} is satisfiable w.r.t. $\mathcal{O}_{\ell, k, \ell', k'}^{\approx}$ iff $\mathcal{D}_{\ell, k}^{\approx}$ is satisfiable w.r.t. \mathcal{O} .

Proof. Assume first that $\mathcal{D}_{\ell, k}^{\approx}$ is satisfiable w.r.t. \mathcal{O} . Then $\mathcal{D}_{\ell, k}^{\approx}$ is satisfiable w.r.t. $\mathcal{O}_{\ell, k, \ell', k'}^{\approx}$ since $\mathcal{O} \models \mathcal{O}_{\ell, k, \ell', k'}^{\approx}$. It follows that \mathcal{D} is satisfiable w.r.t. $\mathcal{O}_{\ell, k, \ell', k'}^{\approx}$, by Lemma 8. Conversely, for an indirect proof, assume that $\mathcal{D}_{\ell, k}^{\approx}$ is not satisfiable w.r.t. \mathcal{O} . By compactness, there is a finite subset \mathcal{F} of $\mathcal{D}_{\ell, k}^{\approx}$ that is not satisfiable w.r.t. \mathcal{O} . We view \mathcal{F} as the head F of a $\ell, k, 1, 1$ -TGD $F \rightarrow \perp$. Then $\mathcal{O} \models F \rightarrow \perp$ and so $F \rightarrow \perp \in \mathcal{O}_{\ell, k, \ell', k'}^{\approx}$. It follows that $\mathcal{D}_{\ell, k}^{\approx}$ is not satisfiable w.r.t. $\mathcal{O}_{\ell, k, \ell', k'}^{\approx}$. But then, by Lemma 8, \mathcal{D} is not satisfiable w.r.t. $\mathcal{O}_{\ell, k, \ell', k'}^{\approx}$. \square

We next describe the careful chase algorithm.

Fix $\ell, k, k' \geq 1$ with $\ell < k$. Let $Q(\bar{x}) = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALCI}, \text{CQ})$, \mathcal{D} a Σ -database, and $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$. We want to decide whether $\bar{a} \in \mathcal{Q}_{\ell, k, 1, k'}^{\approx}(\mathcal{D})$.

We first check whether \mathcal{D} is satisfiable w.r.t. $\mathcal{O}_{\ell, k, 1, k'}^{\approx}$. By Lemma 16, this is the case if $\mathcal{D}_{\ell, k}^{\approx}$ is satisfiable w.r.t. \mathcal{O} . The latter problem can be reduced in polynomial time to the complement of the problem considered in Theorem 5, and thus in EXPTIME and within the time requirements of FPT. Indeed, as we have seen already in Appendix C, $\mathcal{D}_{\ell, k}^{\approx}$ is satisfiable w.r.t. \mathcal{O} iff $\mathcal{O}, \mathcal{D}_{\ell, k}^{\approx} \not\models \exists x A(x)$, where A is a fresh concept name.

We may thus assume that \mathcal{D} is satisfiable w.r.t. $\mathcal{O}_{\ell, k, 1, k'}^{\approx}$. Let $\text{trees}^{1, k'}(q)$ denote the set of all CQs of treewidth $(1, k')$ that can be obtained from q by first quantifying all variables, then taking a contraction, then an induced subquery, and then selecting at most one variable as the answer variable (thus dropping quantification from it). In addition, $\text{trees}^{1, k'}(q)$ contains all AQs $A(x)$ with A a concept name used in \mathcal{O} . Note that all CQs in $\text{trees}^{1, k'}(q)$ are unary or Boolean.

Now extend \mathcal{D} to a database \mathcal{D}' as follows. For every CQ $p(\bar{x}) \in \text{trees}^{1, k'}(q)$ and every $a \in \text{adom}(\mathcal{D})$ with $\langle \mathcal{D}_{\ell, k}^{\approx}, a \rangle, \mathcal{O} \models p(\bar{a})$, take a disjoint copy of $p(x)$ viewed as a database and add it to \mathcal{D} ; if p is unary, then glue the root to the single constant a in \bar{a} . Note that $\langle \mathcal{D}_{\ell, k}^{\approx}, a \rangle, \mathcal{O} \models p(\bar{a})$ is exactly what Theorem 5 allows us to decide, in EXPTIME and within the time requirements of FPT. Note that \mathcal{D}' is homomorphically equivalent to a subdatabase of $\text{ch}_{\mathcal{O}_{\ell, k, 1, k'}^{\approx}}(\mathcal{D})$, which is why we speak of a careful chase. The algorithm now checks whether $\bar{a} \in q(\mathcal{D}')$ using brute force (to attain EXPTIME combined complexity) or using as

a blackbox an algorithm that runs within the time requirements of fixed-parameter tractability (to attain FPT) and returns the result. Correctness is established by the following lemma.

Lemma 17. $\bar{a} \in Q_{\ell,k,1,k'}^{\approx}(\mathcal{D})$ iff $\bar{a} \in q(\mathcal{D}')$.

Proof. We can show the following in exactly the same way as in the proof of Lemma 12.

Claim 1. The restriction of $\text{ch}_{\mathcal{O}_{\ell,k,1,k'}^{\approx}}(\mathcal{D})$ to $\text{adom}(\mathcal{D})$ is identical to the restriction of \mathcal{D}' to $\text{adom}(\mathcal{D})$.

Now for the proof of the lemma. For the ‘if’ direction, assume that $\bar{a} \in Q_{\ell,k,1,k'}^{\approx}(\mathcal{D})$. By Lemma 6, there is a homomorphism h from q to $\text{ch}_{\mathcal{O}_{\ell,k,1,k'}^{\approx}}(\mathcal{D})$ with $h(\bar{x}) = \bar{a}$. It suffices to identify a homomorphism g from q to \mathcal{D}' with $g(\bar{x}) = \bar{a}$.

By Claim 1, we can start the definition of g by setting $g(x) = h(x)$ whenever $h(x) \in \text{adom}(\mathcal{D})$. For completing the definition, let CQ p be obtained from q by quantifying all variables and then contracting by identifying variables x, y whenever $h(x) = h(y)$. Further, let p^- be the restriction of p to those atoms that contain at least one variable x with $h(x) \notin \text{adom}(\mathcal{D})$, viewing the remaining variables x with $h(x) \in \text{adom}(\mathcal{D})$ as answer variables. It is clear that p^- is a collection of CQs of treewidth $(1, k')$ with at most one answer variable each. It suffices to extend g to all these CQs.

Let $\hat{p}(x)$ be a unary CQ in p^- . Then $\mathcal{D}, \mathcal{O}_{\ell,k,1,k'}^{\approx} \models \hat{p}(h(x))$ is witnessed by h . Point 3 of Lemma 8 implies $\langle \mathcal{D}_{\ell,k}^{\approx}, h(x) \rangle, \mathcal{O}_{\ell,k,1,k'}^{\approx} \models \hat{p}(h(x))$. Since $\mathcal{O} \models \mathcal{O}_{\ell,k,1,k'}^{\approx}$, this yields $\langle \mathcal{D}_{\ell,k}^{\approx}, h(x) \rangle, \mathcal{O} \models \hat{p}(h(x))$. By construction of \mathcal{D}' , a copy of \hat{p} has been added to \mathcal{D}' with its root glued to $h(x)$ and thus we can extend g to all variables in $\hat{p}(x)$. The case of Boolean CQs from p^- is analogous. If variables x_1 and x_2 have been identified during the construction of p with the resulting variable called x_1 , then additionally set $g(x_2) = g(x_1)$.

It is readily verified that g is indeed a homomorphism from q to \mathcal{D}' with $g(\bar{x}) = \bar{a}$.

For the ‘only if’ direction, assume that $\bar{a} \in q(\mathcal{D}')$. Then there is a homomorphism h from q to \mathcal{D}' with $h(\bar{x}) = \bar{a}$. It suffices to show that there is a homomorphism g from \mathcal{D}' to $\text{ch}_{\mathcal{O}_{\ell,k,1,k'}^{\approx}}(\mathcal{D})$ with $g(\bar{a}) = \bar{a}$ because then the composition of h with g yields a homomorphism from q to $\text{ch}_{\mathcal{O}_{\ell,k,1,k'}^{\approx}}(\mathcal{D})$ that witnesses that $\bar{a} \in Q_{\ell,k,1,k'}^{\approx}(\mathcal{D})$, as required.

By the claim, we can start by setting $g(c) = h(c)$ for all $c \in \text{adom}(\mathcal{D})$. Now assume that a copy of a unary CQ $p(x) \in \text{trees}(q)$ has been added during the construction of \mathcal{D}' , glueing the root to some $a \in \text{adom}(\mathcal{D})$. Then $\langle \mathcal{D}_{\ell,k}^{\approx}, a \rangle, \mathcal{O} \models p(a)$. Compactness yields a finite subset \mathcal{F} of $\langle \mathcal{D}_{\ell,k}^{\approx}, a \rangle$ with $\mathcal{F}, \mathcal{O} \models p(a)$. We can view \mathcal{F} as a unary CQ $\hat{p}(x)$ of treewidth (ℓ, k) where the free variable x corresponds to the constant a . Then $\mathcal{O} \models \hat{p}(x) \rightarrow p(x)$ and thus $\hat{p}(x) \rightarrow p(x)$ is a CI in $\mathcal{O}_{\ell,k,1,k'}^{\approx}$. Consequently, $\langle \mathcal{D}_{\ell,k}^{\approx}, a \rangle, \mathcal{O}_{\ell,k,1,k'}^{\approx} \models p(a)$ and Point 1 of Lemma 8 yields $\mathcal{D}, \mathcal{O}_{\ell,k,1,k'}^{\approx} \models p(a)$. By Lemma 6, we find a homomorphism h_p from $p(x)$ to $\text{ch}_{\mathcal{O}_{\ell,k,1,k'}^{\approx}}(\mathcal{D})$ with $h_p(x) = a$. We

can extend g to the added copy of $p(x)$ at a based on h_p . The case of Boolean CQs is similar. \square

It remains to analyze the running time of the careful chase algorithm. The initial satisfiability check runs in EXPTIME and FPT and the computation of \mathcal{D}' needs time $2^{O(\|q\|^2)} \cdot |\text{adom}(\mathcal{D})|$ entailment checks, thus overall also running in EXPTIME and FPT. The size of \mathcal{D}' is bounded by $\|\mathcal{D}\| + |\text{adom}(\mathcal{D})| \cdot 2^{O(\|q\|^2)}$. Checking $\bar{a} \in q(\mathcal{D}')$ by brute force means to test $\|\mathcal{D}'\|^{\|q\|}$ candidate homomorphisms, that is, single exponentially many in the case of combined complexity and polynomially many in the case of data complexity, which gives the EXPTIME and PTIME upper bounds in Point 1. If $q \in \text{UCQ}_k^{\text{IW}}$, we can check $\bar{a} \in q(\mathcal{D}')$ using a blackbox procedure that runs in PTIME combined complexity (Dechter and Pearl 1989; Freuder 1990), which yields the FPT upper bound in Point 2.

F Proof of Theorem 5

Theorem 5. Let $\ell, k, k' \geq 1$ with $\ell < k$. Given an OMQ $Q(\bar{x}) = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALC}\mathcal{I}, \text{CQ})$ of arity at most one and with \mathcal{D}_q of treewidth $(1, k')$, a Σ -database \mathcal{D} , and $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$, deciding whether $\bar{a} \in Q(\langle \mathcal{D}_{\ell,k}^{\approx}, \bar{a} \rangle)$ is in EXPTIME in combined complexity and FPT.³

To prove Theorem 5, we use an approach based on two-way alternating parity tree automata. We first introduce this automata model.

Let \mathbb{N} denote the positive integers $\{1, 2, \dots\}$ and let \mathbb{N}^* denote the set of all words over \mathbb{N} viewed as an infinite alphabet. A *tree* is a non-empty (potentially infinite) set $T \subseteq \mathbb{N}^*$ closed under prefixes. The node $\varepsilon \in T$ is the *root* of T . An *infinite path* P in T is a prefix-closed set $P \subseteq T$ such that for every $i \geq 0$, there is a unique word $x \in P$ that is of length i . As a convention, we take $x \cdot 0 = x$ and $(x \cdot c) \cdot -1 = x$. Note that $\varepsilon \cdot -1$ is undefined. We say that T is *m-ary*, $m \geq 1$, if for every $x \in T$, the set $\{i \mid x \cdot i \in T\}$ is of cardinality at most m . W.l.o.g., we assume that all nodes in an m -ary tree are from $\{1, \dots, m\}^*$. For an alphabet Γ and $m \geq 1$, a Γ -labeled *m-ary tree* is a pair (T, L) with T a tree and $L : T \rightarrow \Gamma$ a node labeling function.

For a set X , we use $\mathcal{B}^+(X)$ to denote the set of all positive Boolean formulas over X , i.e., formulas built using conjunction and disjunction over the elements of X used as propositional variables, and where the special formulas `true` and `false` are admitted as well.

Definition 1 (TWAPA). A two-way alternating parity automaton (TWAPA) on m -ary trees is a tuple $\mathfrak{A} = (S, \Gamma, \delta, s_0, c)$ where S is a finite set of states, Γ is a finite alphabet, $\delta : S \times \Gamma \rightarrow \mathcal{B}^+(\text{tran}(\mathfrak{A}))$ is the transition function with $\text{tran}(\mathfrak{A}) = \{\langle i \rangle s, [i] s \mid -1 \leq i \leq m \text{ and } s \in S\}$ the set of transitions of \mathfrak{A} , $s_0 \in S$ is the initial state, and $c : S \rightarrow \mathbb{N}$ is the parity condition that assigns to each state a priority.

Intuitively, a transition $\langle i \rangle s$ with $i > 0$ means that a copy of the automaton in state s is sent to the i -th successor of

³Note that \mathcal{D}_q being of treewidth $(1, k')$ is a stricter condition than q being of treewidth $(1, k')$.

the current node, which is then required to exist. Similarly, $\langle 0 \rangle_s$ means that the automaton stays at the current node and switches to state s , and $\langle -1 \rangle_s$ indicates moving to the predecessor of the current node, which is then required to exist. Transitions $\langle i \rangle_s$ mean that a copy of the automaton in state s is sent to the relevant node if that node exists (which is not required).

Definition 2 (Run, Acceptance). *A run of a TWAPA $\mathfrak{A} = (S, \Gamma, \delta, s_0, c)$ on a Γ -labeled tree (T, L) is a $T \times S$ -labeled tree (T_r, r) such that:*

1. $r(\varepsilon) = (\varepsilon, s_0)$;
2. if $y \in T_r$, $r(y) = (x, s)$, and $\delta(s, L(x)) = \varphi$, then there is a (possibly empty) set $S \subseteq \text{tran}(\mathfrak{A})$ such that S (viewed as a propositional valuation) satisfies φ as well as the following conditions:
 - (a) if $\langle i \rangle_{s'} \in S$, then $x \cdot i$ is defined and there is a node $y \cdot j \in T_r$ such that $r(y \cdot j) = (x \cdot i, s')$;
 - (b) if $\langle i \rangle_{s'} \in S$ and $x \cdot i$ is defined and in T , then there is a node $y \cdot j \in T_r$ such that $r(y \cdot j) = (x \cdot i, s')$.

We say that (T_r, r) is accepting if on all infinite paths in T_r , the maximum priority that appears infinitely often is even. A Γ -labeled tree (T, L) is accepted by \mathfrak{A} if there is an accepting run of \mathfrak{A} on (T, L) . We use $L(\mathfrak{A})$ to denote the set of all Γ -labeled tree accepted by \mathfrak{A} .

It is known (and easy to see) that complement and intersection of TWAPAs can be implemented with only a polynomial blowup. It is also known that their emptiness problem can be solved in time single exponential in the number of states and the maximum priority and polynomial in all other components of the automaton (Vardi 1985). In what follows, we shall generally only explicitly analyze the number of states of a TWAPA, but implicitly also take care that all other components are of the appropriate size for the complexity result that we aim to obtain. In particular, the maximum priority will always be 2.

For the EXPTIME upper bound in Theorem 5, let $\ell, k, k', Q, \mathcal{D}$ and \bar{a} be as in the statement of the theorem with $Q(\bar{x}) = (\mathcal{O}, \Sigma, q)$. We show how to construct a TWAPA \mathfrak{A} with polynomially many states such that $L(\mathfrak{A}) = \emptyset$ iff $\bar{a} \in Q(\langle \mathcal{D}_{\ell, k}^{\approx}, \bar{a} \rangle)$. By what was said above, this yields the desired EXPTIME algorithm.

It is an immediate consequence of the semantics that $\bar{a} \notin Q(\langle \mathcal{D}_{\ell, k}^{\approx}, \bar{a} \rangle)$ if and only if there is model \mathcal{I} of \mathcal{O} and $\langle \mathcal{D}_{\ell, k}^{\approx}, \bar{a} \rangle$ such that $\mathcal{I} \not\models q(\bar{a})$. We call such a model a *witness* for $\bar{a} \notin Q(\langle \mathcal{D}_{\ell, k}^{\approx}, \bar{a} \rangle)$. The witness \mathcal{I} is *forest like* if the interpretation obtained from \mathcal{I} by replacing $r^{\mathcal{I}}$ with $r^{\mathcal{I}} \setminus \text{adom}(\langle \mathcal{D}_{\ell, k}^{\approx}, \bar{a} \rangle) \times \text{adom}(\langle \mathcal{D}_{\ell, k}^{\approx}, \bar{a} \rangle)$ for all role names r is a disjoint union of (potentially infinite) trees. It is well-known that if there is a witness for $\bar{a} \notin Q(\langle \mathcal{D}_{\ell, k}^{\approx}, \bar{a} \rangle)$, then there is also one that is forest like with the tree parts of out-degree at most $|\mathcal{O}|$, see for instance (Lutz 2008b). Witnesses of this kind can be encoded as labeled trees and thus presented as an input to the TWAPA \mathfrak{A} .

We define \mathfrak{A} to run on m -ary Γ -labeled trees where $m = |\mathcal{O}| + (|\text{adom}(\mathcal{D})|^k \cdot 2^k)$ and for $N = \text{adom}(\mathcal{D}) \cup \{1, 2, 3\}$,

the input alphabet is

$$\Gamma = \{\perp\} \cup \{(\mathcal{B}, O) \mid \mathcal{B} \text{ is a } \Sigma \cup \text{sig}(\mathcal{O})\text{-database with} \\ \text{adom}(\mathcal{B}) \subseteq N \text{ and } |\text{adom}(\mathcal{B})| \leq k, \\ \text{and } O \subseteq \text{adom}(\mathcal{B})\}.$$

Informally, \perp is a special label for the root node, the constants from $\text{adom}(\mathcal{D})$ are used to represent the copies of such constants in $\langle \mathcal{D}_{\ell, k}^{\approx}, \bar{a} \rangle$, which is a part of every forest like witness interpretation \mathcal{I} , and the constants $\{1, 2, 3\}$ are used to represent the remaining, tree-shaped parts of \mathcal{I} . If $L(t) = (\mathcal{B}, O)$, then O represents the overlap between $\text{adom}(\mathcal{B})$ and the domain of the database associated with the predecessor of t in the tree, c.f. the definition of ℓ, k -unravelings. For brevity, we then write $L_1(t)$ to denote \mathcal{B} and $L_2(t)$ for O . Whenever we write $L_1(t)$ or $L_2(t)$, we silently assume that $L(t) \neq \perp$. We say that (T, L) is *proper* if it satisfies the following conditions for all $t, t' \in T$:

1. $L(t) = \perp$ if and only if $t = \varepsilon$;
2. if t has predecessor t' , then there is no other child t'' of t' such that $\text{adom}(L_1(t)) = \text{adom}(L_1(t''))$ and $L_2(t) = L_2(t'')$;
3. if t has predecessor $t' \neq \varepsilon$, then $L_2(t') \subseteq \text{adom}(L_1(t))$;
4. if t has predecessor $t' \neq \varepsilon$, then $L_1(t)|_S = L_1(t')|_S$, $S = \text{adom}(L_1(t)) \cap \text{adom}(L_1(t'))$;
5. if $L_1(t) \cap \{1, 2, 3\} \neq \emptyset$, then $|\text{adom}(L_1(t))| \leq 2$ and for all children t' of t , $\text{adom}(L_1(t')) \cap \text{adom}(\mathcal{D}) = \emptyset$.

Formally, every proper Γ -labeled tree (T, L) encodes an interpretation $\mathcal{I}_{(T, L)}$, as follows. First choose a function μ that maps each pair (t, a) with $t \in T \setminus \{\varepsilon\}$ and $a \in \text{adom}(L_1(t))$ to a constant $\mu(t, a)$ such that:

1. if t has predecessor ε , then
 - (a) if q is unary, $\bar{a} = a$, $\text{adom}(L_1(t)) = \{a\}$, and $L_2(t) = \emptyset$, then $\mu(t, a) = a$;⁴
 - (b) otherwise, $\mu(t, a)$ is a fresh constant.
2. if t has predecessor t' and $a \in \text{adom}(L_2(t))$, then $\mu(t, a) = \mu(t', a)$;
3. if $a \notin \text{adom}(L_2(t))$, then $\mu(t, a)$ is a fresh constant.

Then set

$$\begin{aligned} \Delta^{\mathcal{I}_{(T, L)}} &= \{\mu(t, a) \mid t \in T, a \in \text{adom}(L_1(t))\} \\ A^{\mathcal{I}_{(T, L)}} &= \{\mu(t, a) \mid t \in T, A(a) \in L_1(t)\} \\ r^{\mathcal{I}_{(T, L)}} &= \{(\mu(t, a), \mu(t, b)) \mid t \in T, r(a, b) \in L_1(t)\} \end{aligned}$$

for all concept names A and role names r . The *database part* of $\mathcal{I}_{(T, L)}$ is the restriction of $\mathcal{I}_{(T, L)}$ to domain $\{\mu(x, a) \mid x \in T \text{ and } a \in \text{adom}(\mathcal{D})\}$ while the *existential part* is the restriction to domain $\{\mu(x, a) \mid x \in T \text{ and } a \in \{1, 2, 3\}\}$. It is clear by construction that the database part of $\mathcal{I}_{(T, L)}$ has treewidth (ℓ, k) while the existential part has treewidth $(1, 2)$. Moreover, if \mathcal{I} is a witness for $\bar{a} \notin Q(\mathcal{D}_{\ell, k}^{\approx})$ in

⁴Here we assume w.l.o.g. that the copy of $a \in \text{adom}(\mathcal{D})$ in $\mathcal{D}_{\ell, k}^{\approx}$ that the $\langle \cdot \rangle$ operation renames back to a is the copy induced by the ℓ, k -sequence $\{a\}$, c.f. the definition of ℓ, k -unravelings.

which the tree parts have outdegree at most $|\mathcal{O}|$, then we find a Γ -labeled tree (T, L) such that $\mathcal{I}_{(T,L)}$ is isomorphic to \mathcal{I} and the database part of $\mathcal{I}_{(T,L)}$ is identical to the restriction of \mathcal{I} to domain $\text{adom}(\langle \mathcal{D}_{\ell,k}^{\approx}, \bar{a} \rangle)$, up to renaming constants that are not in \bar{a} .

The TWAPA \mathfrak{A} will be obtained as the intersection of four TWAPA, $\mathfrak{A}_{\text{proper}}$, \mathfrak{A}_1 , \mathfrak{A}_2 and \mathfrak{A}_3 . While $\mathfrak{A}_{\text{proper}}$ makes sure that the input tree (T, L) is proper, \mathfrak{A}_1 checks that, up to renaming constants that are not in \bar{a} , $\mathcal{I}_{(T,L)}$ is a model of $\langle \mathcal{D}_{\ell,k}^{\approx}, \bar{a} \rangle$, \mathfrak{A}_2 checks that $\mathcal{I}_{(T,L)}$ is a model of \mathcal{O} , and \mathfrak{A}_3 checks that $\mathcal{I}_{(T,L)} \not\models q(\bar{a})$. In summary, we thus have $L(\mathfrak{A}) \neq \emptyset$ iff there is a witness for $\bar{a} \notin Q(\langle \mathcal{D}_{\ell,k}^{\approx}, \bar{a} \rangle)$. Regarding \mathfrak{A}_1 note that the renaming of constants is unproblematic because $\langle \mathcal{D}_{\ell,k}^{\approx}, \bar{a} \rangle \models Q(\bar{a})$ iff $\mathcal{E} \models Q(\bar{a})$, for all databases \mathcal{E} obtained from $\langle \mathcal{D}_{\ell,k}^{\approx}, \bar{a} \rangle$ by renaming constants that are not in \bar{a} . The TWAPA $\mathfrak{A}_{\text{proper}}$ is very easy to design, we leave details to the reader. A sketch of the construction of the other TWAPAs is provided below.

Details for \mathfrak{A}_1 . This TWAPA has states $s_{|I}$ and $s_{I,O}$ for every $I \subseteq \text{adom}(\mathcal{D})$ with $1 \leq |I| \leq k$ and every $O \subseteq I$. When in state s_I or $s_{I,O}$, the TWAPA verifies that the current symbol encodes an interpretation that is a model of $\mathcal{D}_{|I}$. There is an additional starting state s_0 from which the TWAPA transitions to every child, making sure that for every $I \subseteq \text{adom}(\mathcal{D})$ with $1 \leq |I| \leq k$, some child is visited in state s_I . From state s_I or $s_{I,O}$, the TWAPA again transitions to every child, making sure that for every $J \subseteq \text{adom}(\mathcal{D})$ with $1 \leq |J| \leq k$ and every $O \subseteq J$, some child is visited in state $s_{J,O}$. The number of states needed for this TWAPA is $O(|\text{adom}(\mathcal{D})|^k \cdot 2^k)$ and thus polynomial since k is fixed.

Details for \mathfrak{A}_2 . We can assume \mathcal{O} to be of the form $\{\top \sqsubseteq C_{\mathcal{O}}\}$, where $C_{\mathcal{O}}$ is in negation normal form, that is, negation is applied only to concept names, but not to compound concepts. The TWAPA visits all nodes x in the input tree, considers every constant a in $L_1(x)$, and verifies that $\mu(x, a)$ satisfies $C_{\mathcal{O}}$ in $\mathcal{I}_{(T,L)}$. This can be done using states of the form $s_{C,a}$, with $C \in \text{sub}(C_{\mathcal{O}})$ and $a \in N$, where the purpose of $s_{C,a}$ is to verify that $\mu(x, a) \in C^{\mathcal{I}_{(T,L)}}$ for the current node x . When in state $s_{C_1 \sqcap C_2, a}$, for instance, the automaton stays at the same node and transitions to both states $s_{C_1, a}$ and $s_{C_2, a}$. The only non-trivial cases are concepts of the form $\exists r.C$ and $\forall r.C$ since r -successors of $\mu(x, a)$ in $\mathcal{I}_{(T,L)}$ can be generated by constants in some $L_1(y)$ with $y \neq x$. In the $\forall r.C$ case, the automaton considers all r -successors of a in $L_1(x)$ and then moves to all children y of x with $a \in L_2(y)$ and to the parent y of x if $a \in L_2(x)$ and considers all r -successors of a in $L_1(y)$. Concepts $\exists r.C$ are dealt with in a dual way. Instead of sending copies of itself to all neighbors, the TWAPA guesses a neighbor to send a copy of itself to. The parity condition is used to ensure that the TWAPA eventually arrives at a node in which $\exists r.C$ is satisfied for a . The number of states is $O(|\mathcal{O}| \cdot k)$.

Details for \mathfrak{A}_3 . Since TWAPAs can be complemented without increasing the number of states, it suffices to construct a TWAPA that checks that $\mathcal{I}_{(T,L)} \models q(\bar{a})$. We characterize the existence of a homomorphism h from q to $\mathcal{I}_{(T,L)}$ with

$h(\bar{x}) = \bar{a}$ by the existence of a match witness, as defined below.

Let (T, L) be a proper Γ -labeled tree. We say that nodes $t_1, t_2 \in T$ are a -connected, a a constant from N , if there is a path s_1, \dots, s_n in T with $t_1 = s_1, t_2 = s_n$, and such that for $1 \leq i < n$, either s_{i+1} is a successor of s_i and $a \in L_2(s_{i+1})$ or s_{i+1} is the predecessor of s_i and $a \in L_2(s_i)$. Fix a $(1, k')$ -tree decomposition $T_q = (V_q, E_q, B_q)$ of q . A *match witness* for q in (T, L) is a triple (D, P, H) where

- D assigns to each $v \in V_q$ a tree decomposition $D(v) = (V_v, E_v, B_v)$ of $q|_{B_q(v)}$;
- P assigns to each $v \in V_q$ and $\hat{v} \in V_v$ a node $P(v, \hat{v})$ from the input tree T ;
- H assigns to each $v \in V_q$ and $\hat{v} \in V_v$ a homomorphism $H(v, \hat{v})$ from $q|_{B_q(\hat{v})}$ to $L_1(P(v, \hat{v}))$.

such that the following conditions are satisfied:

1. for all $u, v \in V_q, \hat{u} \in V_u, \hat{v} \in V_v$, and $x \in B_u(\hat{u}) \cap B_v(\hat{v})$, $H(u, \hat{u})(x) = H(v, \hat{v})(x)$ and $P(u, \hat{u})$ and $P(v, \hat{v})$ are $H(u, \hat{u})(x)$ -connected in T ;
2. if q is unary with answer variable x and $\bar{a} = a$, then for all $v \in V_q$ and $\hat{v} \in V_v$ such that $x \in B_v(\hat{v})$, $H(v, \hat{v})(x) = a$ and $P(v, \hat{v})$ is a -connected to the unique successor t of \hat{v} with $\text{adom}(L_1(t)) = \{a\}$ and $L_2(t) = \emptyset$ (please see Condition 1(a) in the definition of $\mu(t, a)$).

Lemma 18. *Let (T, L) be a proper Γ -labeled tree. There is a match witness for q in (T, L) if and only if $\mathcal{I}_{(T,L)} \models q(\bar{a})$.*

Proof. “ \Rightarrow ”. Let (D, P, H) be a match witness for q in (T, L) . To show that $\mathcal{I}_{(T,L)} \models q(\bar{a})$, we construct a homomorphism h from q to $\mathcal{I}_{(T,L)}$ with $h(\bar{x}) = \bar{a}$. Let x be a variable in q . Then there is a $v \in V_q$ and a $\hat{v} \in V_v$ such that $x \in B_v(\hat{v})$. Set $h(x)$ to be the element of $\Delta^{\mathcal{I}_{(T,L)}}$ generated by $H(v, \hat{v})(x)$, that is, $\mu(P(v, \hat{v}), H(v, \hat{v})(x))$. Due to the connectedness condition of tree decompositions and by Condition 1, the choices of v and \hat{v} do not matter and h is well-defined. Moreover, h is clearly a homomorphism from q to $\mathcal{I}_{(T,L)}$ since each $H(v, \hat{v})$ is a homomorphism from $q|_{B_q(\hat{v})}$ to $L_1(P(v, \hat{v}))$. Finally $h(\bar{x}) = \bar{a}$ due to Condition 2.

“ \Leftarrow ”. Let $\mathcal{I}_{(T,L)} \models q(\bar{a})$ and let h be a homomorphism from q to $\mathcal{I}_{(T,L)}$ with $h(\bar{x}) = \bar{a}$. We define the match witness (D, P, H) as follows. Let $v \in V_q$. For every node t in the input tree, let $W_{v,t}$ denote the set of all variables $x \in B_q(v)$ such that for some constant a that occurs in $L_1(t)$, $h(x)$ is the element of $\Delta^{\mathcal{I}_{(T,L)}}$ generated by a in t , that is, $h(x) = \mu(t, a)$ where μ is the function from the definition of $\mathcal{I}_{(T,L)}$.

- V_v is the \subseteq -minimal prefix-closed subset of T such that for all $x, y \in B_q(v)$ that cooccur in an atom in q , there is a $t \in V_v$ with $\{x, y\} \subseteq W_{v,t}$.
- $E_v = \{(t, t \cdot i) \mid t \cdot i \in V_v, i \in \{1, \dots, m\}\}$;
- $B_v(\hat{v}) = W_{v, \hat{v}}$;
- $P(v, \hat{v}) := \hat{v}$, and
- $H(v, \hat{v})(x)$ is defined to be the constant that generates $h(x)$, that is, $\mu(\hat{v}, x)$, for all $x \in W_{v, \hat{v}}$.

It follows from the definition of $\mathcal{I}_{(T,L)}$ that (V_v, E_v, B_v) is a tree decomposition of $q|_{B_q(v)}$. Using that h is a homomorphism and Condition 3 of proper input trees, it can further be shown that $H(v, \hat{v})$ is a homomorphism from $q|_{B_v(\hat{v})}$ to $L_1(P(v, \hat{v}))$, as required. Moreover, Conditions 1 and 2 of match witnesses are clearly satisfied. \square

Note that the proof of Lemma 18 does not rely on the fact that q has treewidth $(1, k')$, so the lemma holds even if T_q is an (ℓ', k') -tree decomposition with $\ell' > 1$. However, we aim to show that the existence of a match tree can be checked using a TWAPA \mathfrak{A}_3 with polynomially many states, and this is where the assumption $\ell' = 1$ is needed.

The general idea is that \mathfrak{A}_3 processes the bags of T_q in a top-down manner, always storing the current node $v \in V_q$ in its state. For every $v \in V_q$, it verifies the existence of a tree decomposition $D(v)$ of $q|_{B_q(v)}$, also in a top-down manner. Note, however, that $D(v)$ is not fixed like T_q , but rather the automaton needs to construct it, that is, we must be able to extract the decompositions $D(v)$ from a successful run. While verifying that the required $D(v)$ exists, the automaton also verifies the existence of suitable P and H components of the match witness. When considering a node \hat{v} of decomposition $D(v)$, it moves to the node in the input tree that is $P(v, \hat{v})$ and then checks that the required homomorphism $H(v, \hat{v})$ exists.

We use states of the form $s_{v,M,h}$, where

- $v \in V_q$ is the current node of T_q ,
- $M : B_q(v) \rightarrow \{\text{here}, \text{below}, \text{else}\}$ records which variables from $B_q(v)$ are part of the bag of $D(v)$ that is currently being treated ('here'), which are not in the current bag, but in a bag strictly below ('below'), and which are only part of bags that are neither the current one nor below it ('else'), and
- h is a homomorphism from $q|_{\{x \in B_q(v) \mid M(x) = \text{here}\}}$ to the database $L_1(t)$ associated with the current node t of the input tree.

Note that the number of variables in $B_q(v)$ is bounded by k' , and thus the number of possible functions M is polynomial. Including some book keeping states, in fact, the number of required states is $O(|q|^2 \cdot 3^{k'} \cdot k^{k'})$.

\mathfrak{A}_3 send a copy of itself to some node of the input tree and then starts in state $s_{v,M,h}$ where v is the root of T_q , h the empty homomorphism, and M maps every variable to here or below. Now assume it is in some state $s_{v,M,h}$ which (implicitly) represents some node \hat{v} of $D(v)$. It then spawns $n \leq |B_q(v)|$ copies of itself, corresponding to children $\hat{v}_1, \dots, \hat{v}_n$ of \hat{v} in $D(v)$, in states $s_{v,M_1,h_1}, \dots, s_{v,M_n,h_n}$. In the M_i , each variable that was here before is now here or else and each variable that was else before is still else. Moreover, each variable x that was below before is now here or below or else, but is here or below in exactly one M_i . What is more, h_i agrees with h on the intersection of their domains. The latter partly achieves Condition 1 from the definition of match witnesses. The automaton then moves each of the n copies to a node in the input tree, corresponding to the choice of $P(\hat{v}_i)$. To make sure that it eventually reaches that node, we use the parity condition. More-

over, we make sure that for every variable x that is in the domain of both h and h_i , the path is $h(x)$ -connected. This fully achieves Condition 1.

The TWAPA might need to spawn additional copies to transition to successors of v in the 'outer' tree decomposition T_q . For successors u of v such that $B_q(u) \cap B_q(v) = \emptyset$, this is done as soon as the processing of v starts. For successors u of v such that $B_q(u) \cap B_q(v) = \{x\}$, this is done the first time that $M(x)$ is set to here. If the current state is $s_{v,h,M}$, we then spawn the new copy in a state $s_{u,h',M'}$, where h' is the homomorphism that maps only x to $h(x)$ and M' sets x to here and all other variables from $B_q(u)$ to below. We thus start the top-down verification of $D(u)$ at a bag that contains the node x that $B_q(u)$ shares with its predecessor bag $B_q(v)$ in T_q . Note that this part fails for $\ell' > 1$ since then $B_q(u)$ and $B_q(v)$ can share two variables x_1, x_2 and there is no guarantee that there is a bag of $D(v)$ in which both x_1 and x_2 are mapped to here.

Finally, we also need to verify Condition 2 from the definition of match witnesses. This is easy. If q has answer variable x and $M(x) = \text{here}$, then \mathfrak{A}_3 makes sure that $h(x) = a$ (with $\bar{a} = a$) and it spawns a copy that verifies that the current node is \bar{a} -connected to the (unique) successor t of ε with $\text{adom}(L_1(t)) = \{a\}$ and $L_2(t) = \emptyset$.

G Proofs for Section 7

Lemma 5. *Let $Q = (\mathcal{O}, \Sigma, q) \in (FO, UCQ)$ and $\ell, k \geq 1$ with $\ell < k$. Then for all Σ -databases \mathcal{D} and $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$, $\bar{a} \in \text{app}_{\mathcal{D}, \varepsilon, k}(Q, \mathcal{D})$ iff $\bar{a} \in Q(\mathcal{D}_{\bar{a}, \ell, k}^{\approx})$.*

Proof. We generalize the proof of Lemma 2. For the " \supseteq " direction, observe that for every finite subset \mathcal{D}' of $\mathcal{D}_{\bar{a}, \ell, k}^{\approx}$ containing \bar{a} the pointed database (\mathcal{D}', \bar{a}) is a member of $\mathcal{D}_{\ell, k}$. We also have a homomorphism from \mathcal{D}' to \mathcal{D} mapping \bar{a} to \bar{a} . Thus, this direction is again a consequence of compactness.

For the " \subseteq " direction, assume that there is a pointed database (\mathcal{D}', \bar{b}) of treewidth (ℓ, k) , $\bar{b} \in Q(\mathcal{D}')$, and there is a homomorphism h from \mathcal{D}' to \mathcal{D} such that $h(\bar{b}) = \bar{a}$. To show that $\bar{a} \in Q(\mathcal{D}_{\bar{a}, \ell, k}^{\approx})$, it suffices to construct a homomorphism g from \mathcal{D}' to $\mathcal{D}_{\bar{a}, \ell, k}^{\approx}$ with $g(\bar{b}) = \bar{a}$ and to recall that answers to OMQs are preserved under taking homomorphic images of databases. But the existence of such a homomorphism is a consequence of Lemma 7. \square

Proposition 4. *Let $\ell, k, \ell', k' \geq 1$ with $\ell < k$ and $\ell' < k'$.*

1. *In $(\mathcal{ALCI}, UCQ_{\ell, k}^{\text{tw}})$, ℓ, k -database relaxing OMQ evaluation is at least as complete as ℓ, k, ℓ', k' -ontology relaxing OMQ evaluation.*
2. *For OMQ from $(\mathcal{ALCI}, CQ_{\ell', k'}^{\text{tw}})$ of arity at most $r \leq 1$, $\ell + r, k + r, \ell', k'$ -ontology relaxing OMQ evaluation is at least as complete as ℓ, k -database relaxing OMQ evaluation.*

Proof. For the proof of Point 1, assume that Q is an OMQ in $(\mathcal{ALCI}, UCQ_{\ell, k}^{\text{tw}})$. Assume that $\bar{a} \in Q_{\ell, k, \ell', k'}^{\approx}(\mathcal{D})$. Then, by Lemma 9, $\bar{a} \in Q_{\ell, k, \ell', k'}^{\approx}(\mathcal{D})$ iff $\bar{a} \in Q_{\ell, k, \ell', k'}^{\approx}(\mathcal{D}_{\bar{a}, \ell, k}^{\approx})$. Thus, $\bar{a} \in Q(\mathcal{D}_{\bar{a}, \ell, k}^{\approx})$ since $\mathcal{O} \models \mathcal{O}_{\ell, k, \ell', k'}^{\approx}$.

For the proof of Point 2, assume that $Q = (\mathcal{O}, \Sigma, q(\bar{x}))$ is an OMQ in $(\mathcal{ALCCl}, \text{CQ}_{\ell', k'}^{\text{tw}})$ of arity at most $r \leq 1$. Let $\bar{a} \in Q(\mathcal{D}_{\bar{a}, \ell, k}^{\approx})$. Compactness yields a finite subset \mathcal{F} of $\mathcal{D}_{\bar{a}, \ell, k}^{\approx}$ with $\mathcal{F}, \mathcal{O} \models q(\bar{a})$. We can view \mathcal{F} as a CQ $p(\bar{x})$ of treewidth $(\ell + r, k + r)$ where \bar{x} corresponds to \bar{a} . Then $\mathcal{O} \models p(\bar{x}) \rightarrow q(\bar{x})$ and thus $p(\bar{x}) \rightarrow q(\bar{x})$ is a TGD in $\mathcal{O}_{\ell+r, k+r, \ell', k'}^{\approx}$. Consequently, $\bar{a} \in Q_{\ell+r, k+r, \ell', k'}^{\approx}(\mathcal{D})$, as required. \square

We prove the upper bounds in Theorem 6.

Theorem 6. *Let $\ell, k \geq 1$, $\ell < k$. Then ℓ, k -database relaxing OMQ evaluation is*

1. 2EXPTIME-complete in combined complexity and fixed-parameter tractable with double exponential running time in $(\mathcal{ALCCl}, \mathcal{Q})$, $\mathcal{Q} \in \{\text{CQ}, \text{UCQ}, \text{CQ}_p^{\text{tw}}, \text{UCQ}_p^{\text{tw}} \mid p \geq 1\}$;
2. EXPTIME-complete in combined complexity and fixed-parameter tractable with single exponential running time in $(\mathcal{ALCCl}, \mathcal{Q})$ and $(\mathcal{ALCCl}, \mathcal{Q})$, $\mathcal{Q} \in \{\text{AQ}, \text{BELIQ}\}$.

We prove the upper bounds in Theorem 6 using an elimination algorithm that is inspired by algorithms for achieving k -consistency of constraint satisfaction problems. It also bears similarity to type elimination in modal logic. We start with Point 2 and then proceed to Point 1. For Point 2, we consider a slight generalization in preparation of the proof of Point 1 later on. By Lemma 5, deciding $\bar{a} \in \text{app}_{\mathcal{D}_{\ell, k}}(\mathcal{D})$ amounts to deciding $\bar{a} \in Q(\mathcal{D}_{\bar{a}, \ell, k}^{\approx})$. Based on this, the announced generalization of Point 2 is as follows.

Theorem 10. *Fix $\ell, k \geq 1$ with $\ell < k$. Given an OMQ $Q(\bar{x}) = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALCCl}, \text{bELIQ})$, a Σ -database \mathcal{D} , $S \subseteq \text{adom}(\mathcal{D})$, and $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$, it is fixed-parameter tractable with single exponential running time to decide whether $\bar{a} \in Q(\mathcal{D}_{S, \ell, k}^{\approx})$.*

To prove Theorem 10, fix $\ell, k \geq 1$ with $\ell < k$. Let $Q(\bar{x}) = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALCCl}, \text{bELIQ})$, \mathcal{D} be a Σ -database, $S \subseteq \text{adom}(\mathcal{D})$, and $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$. Before we formulate the algorithm, we introduce several notions that it uses. We use $\text{c1}(\mathcal{O}, q)$ to denote the smallest set that contains all concepts in \mathcal{O} and the bELIQ q , viewed as an \mathcal{ELT}^u -concept, and that is closed under subconcepts and single negation.

For a set $L \subseteq \text{adom}(\mathcal{D})$ with $S \subseteq L$, an L -assignment μ is a set that contains exactly one of $C(a)$, $\neg C(a)$ for all $\neg C \in \text{cl}(\mathcal{O}, q)$ and $a \in L$, subject to the following conditions:

1. if q is an ELIQ and $\bar{a} = a_0$, then $\neg q(a_0) \in L$;
2. if q is a BELIQ, then $\neg q(a) \in L$ for all $a \in L$.

We shall be interested in L -assignments with $L \subseteq \text{adom}(\mathcal{D})$, $S \subseteq L$, and $|L \setminus S| \leq \ell$. When speaking of sets $L \subseteq \text{adom}(\mathcal{D})$, we from now on silently assume that L satisfies the latter two conditions.

The L -assignment $\text{as}_{\mathcal{I}}(L)$ realized in a model \mathcal{I} of \mathcal{O} with $L \subseteq \Delta^{\mathcal{I}}$ is the set that contains $C(a)$ for all $C \in \text{cl}(\mathcal{O}, q)$ and $a \in L$ with $a \in C^{\mathcal{I}}$.

A k -subset of $\text{adom}(\mathcal{D})$ is a set $K \subseteq \text{adom}(\mathcal{D})$ such that $S \subseteq K$ and $|K \setminus S| \leq k$. Let Γ be a set of L -assignments and let K be a k -subset. A Γ -choice for K is a subset $\Gamma' \subseteq \Gamma$

that contains exactly one L -assignment for each $L \subseteq K$. We say that Γ' is *consistent* if the extended database

$$\mathcal{D}|_K \cup \{C(a) \mid \exists \mu \in \Gamma' : C(a) \in \mu\}$$

is satisfiable with the \mathcal{ALCCl} -ontology $\mathcal{O} \cup \{\top \sqsubseteq C\}$ where $C = D$ if $q = \exists u.D(x)$ is a BELIQ and $C = \top$ otherwise.

We are now ready to describe the algorithm. It computes a sequence $\Gamma_0, \Gamma_1, \dots$ of sets of L -assignments, starting with the set Γ_0 of all L -assignments for any $L \subseteq K$. When constructing Γ_{i+1} from Γ_i , it considers all L -assignments $\mu \in \Gamma_i$ and checks whether for all k -subsets $K \supseteq L$, there is a consistent Γ_i -choice Γ for K such that $\mu \in \Gamma$. If this is the case, then μ remains in Γ_{i+1} . Otherwise, it is removed. Let $f \geq 0$ be smallest such that $\Gamma_f = \Gamma_{f+1}$. The algorithm answers ‘yes’ if Γ_f is empty and ‘no’ otherwise.

Before we prove correctness, let us analyze the running time of the algorithm. We start with the number of iterations. The number of sets $L \subseteq \text{adom}(\mathcal{D})$ with $|L| \leq \ell$ is bounded by $(\|\mathcal{D}\| + 1)^\ell$, thus polynomial in $\|\mathcal{D}\|$. For each such L , the number of L -assignments is in $2^{O(\|\mathcal{Q}\|^2)}$. Since at least one L -assignment is eliminated in each iteration until we reach Γ_f , the number of iterations is bounded by $2^{O(\|\mathcal{Q}\|^2)} \cdot (\|\mathcal{D}\| + 1)^\ell$. It remains to consider a single iteration. The algorithm goes through all sets L and all L -assignments μ , thus considers $2^{O(\|\mathcal{Q}\|^2)} \cdot (\|\mathcal{D}\| + 1)^\ell$ assignments. It then considers all k -subsets $K \supseteq L$ of which there are at most $(\|\mathcal{D}\| + 1)^k$ many. Next, it goes through all Γ_i -choices for K . As there are only $(k + 1)^\ell$ sets $L \subseteq K$ with $|L| \leq \ell$, the number of such choices is also bounded single exponentially in $\|\mathcal{O}\| + \|\mathcal{Q}\|$ (with $(k + 1)^\ell$ appearing as a constant in the exponent). Overall, we achieve fixed-parameter tractability with single exponential running time, as desired.

Lemma 19. *The algorithm answers ‘yes’ iff $\mathcal{D}_{S, \ell, k}^{\approx} \models Q(\bar{a})$.*

Proof. For the ‘only if’ direction, assume $\mathcal{D}_{S, \ell, k}^{\approx} \not\models Q(\bar{a})$. We need to show that the algorithm answers ‘no’. Since $\mathcal{D}_{S, \ell, k}^{\approx} \not\models Q(\bar{a})$, there is a model \mathcal{I} of \mathcal{O} and $\mathcal{D}_{S, \ell, k}^{\approx}$ with $\mathcal{I} \not\models q(\bar{a})$. For every $b \in \text{adom}(\mathcal{D}_{S, \ell, k}^{\approx})$, define $\text{tp}(b) = \{C \in \text{c1}(\mathcal{O}, q) \mid b \in C^{\mathcal{I}}\}$. For every ℓ, k -sequence $v = S_0, O_0, \dots, O_{n-1}, S_n$ with $n \geq 1$ we define $L_v \subseteq \text{adom}(\mathcal{D})$ as $h(O_{n-1})$, where $h : \text{adom}(\mathcal{D}_{S, \ell, k}^{\approx}) \rightarrow \text{adom}(\mathcal{D})$ is the uncopying map, and we define the L_v -assignment $\mu_v = \{C(a) \mid C \in \text{tp}(a_v), a \in L_v\}$. Since \mathcal{I} is a model of \mathcal{O} and $\mathcal{D}_{S, \ell, k}^{\approx}$ with $\mathcal{I} \not\models q(\bar{a})$, μ_v is indeed satisfies Conditions 1 and 2 of L_v -assignments. Let $\Psi = \{\mu_v \mid v \text{ is an } \ell, k\text{-sequence of length } n \geq 1\}$. We argue by induction on i that $\Psi \subseteq \Gamma_i$ for every $i \geq 0$. Since Γ_0 is the set of all L -assignments, $\Psi \subseteq \Gamma_0$. For the induction step, assume $\Psi \subseteq \Gamma_i$ for some $i \geq 0$. Let $\mu_v \in \Psi$ be an L_v -assignment for some $v = S_0, O_0, \dots, O_{n-1}, S_n$. We need to show that for every k -subset K with $K \supseteq L_v$ there is a consistent Γ_i -choice Γ for K such that $\mu_v \in \Gamma$. Let K be a k -subset. For every $L \subseteq K$, choose a set $S_{n+1} \supseteq L$ with $|S_{n+1} \setminus S| \leq k$ and consider the sequence $v_{K, L} = S_0, O_0, \dots, O_{n-1}, K, L, S_{n+1}$. This is an ℓ, k -sequence that yields an L -assignment $\mu_{v_{K, L}} \in \Psi$. The set of all such L -assignments forms a Γ_i -choice Γ . Since \mathcal{I}

is a model that realizes all L -assignments from Γ , Γ is consistent, so μ_v remains in Γ_{i+1} . Since Ψ is not empty, this implies that the algorithm answers ‘no’.

For the ‘if’ direction, assume the algorithm answers ‘no’. We need to show that $\mathcal{D}_{S,\ell,k}^{\approx} \not\models Q(\bar{a})$. Since the algorithm answers ‘no’, the set $\Gamma := \Gamma_f$ is not empty. This implies that Γ contains at least one L -assignment for every set $L \subseteq \text{adom}(\mathcal{D})$ with $S \subseteq L$ and $|L \setminus S| \leq \ell$, since as soon as there is an L_0 such that some Γ_i does not contain an L_0 -assignment, the algorithm eventually eliminates all L -assignments and thus returns ‘yes’. Thus, for every L -assignment $\mu \in \Gamma$ and every k -subset K , there is a consistent Γ -choice $\Gamma_{\mu,K}$ for K such that $\mu \in \Gamma_{\mu,K}$. We construct a model \mathcal{I} of $\mathcal{D}_{S,\ell,k}^{\approx}$ and \mathcal{O} with $\mathcal{I} \not\models q(\bar{a})$. First, we define a function $t : \text{adom}(\mathcal{D}_{S,\ell,k}^{\approx}) \rightarrow 2^{\text{cl}(\mathcal{O},q)}$ that assigns a type $t(a)$ to every $a \in \text{adom}(\mathcal{D}_{S,\ell,k}^{\approx})$ and in a second step, we construct the model \mathcal{I} using t .

We start by defining t on S . Choose any S -assignment $\mu_0 \in \Gamma$ and define $t(a) = \{C \in \text{cl}(\mathcal{O},q) \mid C(a) \in \mu_0\}$ for every $a \in S$. We proceed to define t inductively on the length of the ℓ, k -sequence in which the element appears first. For sequences of length 0, let $v = S_0$ be an ℓ, k -sequence. For every $a_v \in \text{adom}(\mathcal{D}_{S,\ell,k}^{\approx})$ that is a copy of $a \in S_0 \setminus S$, define $t(a_v)$ according to Γ_{μ_0,S_0} , i.e. choose any $\mu \in \Gamma_{\mu_0,S_0}$ that is defined on a and set $t(a_v) = \{C \in \text{cl}(\mathcal{O},q) \mid C(a) \in \mu\}$. If t has already been defined for all elements that appear in ℓ, k -sequences of length at most i , let $v = S_0, O_0, S_1, \dots, S_i, O_i, S_{i+1}$. Since $O_i \subseteq S_i$, t has already been defined (by means of an O_i -assignment μ) on elements of $\text{adom}(\mathcal{D}_{S,\ell,k}^{\approx})$ that are copies of elements of O_i . For every $a_v \in \text{adom}(\mathcal{D}_{S,\ell,k}^{\approx})$ that is a copy of an $a \in S_{i+1} \setminus S_i$ define $t(a_v)$ according to $\Gamma_{\mu,S_{i+1}}$, i.e. choose any $\mu' \in \Gamma_{\mu,S_{i+1}}$ that is defined on a and set $t(a_v) = \{C \in \text{cl}(\mathcal{O},q) \mid C(a) \in \mu'\}$. This finishes the definition of t .

Now we construct the model \mathcal{I} of $\mathcal{D}_{S,\ell,k}^{\approx}$ and \mathcal{O} with $\mathcal{I} \not\models q(\bar{a})$. For every $a \in \text{adom}(\mathcal{D}_{S,\ell,k}^{\approx})$, choose a tree-shaped model \mathcal{I}_a of $t(a)$. The interpretation \mathcal{I} is defined as follows: Start from $\mathcal{D}_{S,\ell,k}^{\approx}$ seen as an interpretation and for every $a \in \text{adom}(\mathcal{D}_{S,\ell,k}^{\approx})$, add a copy of \mathcal{I}_a and identify a with the root of \mathcal{I}_a . It can be verified that \mathcal{I} is indeed a model of \mathcal{O} with $\mathcal{I} \not\models q(\bar{a})$. \square

We now turn to Point 1 of Theorem 6. Fix $\ell, k \geq 1$ with $\ell < k$. Let $Q(\bar{x}) = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALCCI}, \text{UCQ})$, \mathcal{D} be a Σ -database, and $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$. By Lemma 5, it suffices to decide whether $\bar{a} \in Q(\mathcal{D}_{\bar{a},\ell,k}^{\approx})$. To reduce complexity, we first replace Q with a set of Boolean OMQs, similar to what was done in the proof of Theorem 2.

We decorate \mathcal{D} and decorate as well as simplify q . Let $\bar{x} = x_0 \cdots x_{n-1}$ and $\bar{a} = a_0 \cdots a_{n-1}$. For each (possibly inverted) role name r used in q and all $i < n$, introduce a fresh concept name A_{r,a_i} that, informally, expresses the existence of an r -edge to a_i . Extend \mathcal{D} to a database \mathcal{D}^d by adding $A_{r,a_i}(b)$ for all $r(b, a_i) \in \mathcal{D}$. This is clearly possible in time $O(\|\mathcal{D}\| \cdot \|q\|^2)$. Further, let q^d be obtained from q^c

by doing the following for each CQ p in q^c :

1. replace each atom $r(y, a_i)$, y a quantified variable, with $A_{r,a_i}(y)$;
2. for each atom $r(x_{i_1}, x_{i_2})$, $0 \leq i_1, i_2 < n$, check whether $r(a_{i_1}, a_{i_2}) \in \mathcal{D}$; if the check fails, remove p from the UCQ; if it succeeds, remove atom $r(x_{i_1}, x_{i_2})$ from p ;
3. for each atom $A(x_i)$, $0 \leq i < n$, check whether $\mathcal{D}_{\bar{a},\ell,k}^{\approx}, \mathcal{O} \models A(a_i)$ using Theorem 10 with $S = \bar{a}$; if the check fails, remove p from the UCQ; if it succeeds, remove atom $A(x_i)$ from p .

If some of the CQs in q^c becomes empty in the process (that is, all of its atoms have been dropped), then we know that $\bar{a} \in Q^c(\mathcal{D}_{\bar{a},\ell,k}^{\approx})$ and return ‘yes’. If all CQs have been dropped, then we know that $\bar{a} \notin Q^c(\mathcal{D}_{\bar{a},\ell,k}^{\approx})$ and return ‘no’. Clearly all of the above can be done in time $2^{O(\|q\|^2)} \cdot O(\|\mathcal{D}\|)$. Note that q^d contains no answer variables as all atoms that mention them have been dropped. It is not hard to show the following using some basic manipulations of homomorphisms that witness query answers.

Lemma 20. $\bar{a} \in Q^c(\mathcal{D}_{\bar{a},\ell,k}^{\approx})$ iff $(\mathcal{D}^d)_{\bar{a},\ell,k}^{\approx}, \mathcal{O} \models q^d$.

It thus remains to decide whether $(\mathcal{D}^d)_{\bar{a},\ell,k}^{\approx}, \mathcal{O} \models q^d$. We may view q^d as a Boolean UCQ in which each CQ is a disjunction of such CQs, that is, a disjunction of conjunctions of such CQs. Using the laws of distributivity, we can convert q^d into an equivalent conjunction of disjunctions of CQs $q_1 \wedge \cdots \wedge q_k$. To decide whether $(\mathcal{D}^d)_{\bar{a},\ell,k}^{\approx}, \mathcal{O} \models q^d$, it suffices to decide whether $(\mathcal{D}^d)_{\bar{a},\ell,k}^{\approx}, \mathcal{O} \models q_i$ for $1 \leq i \leq k$. In the following, we concentrate on this task.

For the sake of readability, let us summarize the remaining task and reorganize notation. We still have fixed $\ell, k \geq 1$ with $\ell < k$. We further have a Boolean OMQ $Q = (\mathcal{O}, \Sigma, q) \in (\mathcal{ALCCI}, \text{UCQ})$ with each CQ in q connected and a Σ -database \mathcal{D} . We want to decide whether $\mathcal{D}_{\bar{a},\ell,k}^{\approx} \models Q$.

We reuse and suitable adapt the notions and algorithm from the proof of Theorem 10. With $\text{cl}(q)$, we denote the set of all CQs that can be obtained from a CQ in q as follows:

1. take a contraction of treewidth (ℓ, k) ;
2. take a subquery, that is, the atoms in the new CQ are a (non-empty) subset of those in the given one;
3. choose zero or more quantified variables and make them answer variables.

We shall consider CQs that we obtain from a CQ $p(\bar{y})$ in $\text{cl}(q)$ by replacing the answer variables \bar{y} with a tuple of constants \bar{c} . We denote the result of such a replacement with $p(\bar{c})$ and call it an *instantiation* of p .

We use $\text{cl}(\mathcal{O}, q)$ to denote the smallest set that contains all concepts in \mathcal{O} and all ELIQs in $\text{cl}(q)$, viewed as \mathcal{ELI} -concepts, and that is closed under subconcepts and single negation.

For a set $L \subseteq \text{adom}(\mathcal{D})$ with $S \subseteq L$, an L -assignment μ is a set that contains

- exactly one of $C(a), \neg C(a)$ for all $\neg C \in \text{cl}(\mathcal{O}, q)$ and $a \in L$, and

- exactly one of $p(\bar{c}), \neg p(\bar{c})$ for all instantiations $p(\bar{c})$ of a CQ from $\text{cl}(q)$ with $\bar{c} \subseteq L$.

We require that μ does not contain any contractions of q (thus it contains the negation of all such contractions of treewidth (ℓ, k)). We shall be interested in L -assignments with $L \subseteq \text{adom}(\mathcal{D}), S \subseteq L$, and $|L \setminus S| \leq \ell$. When speaking of sets $L \subseteq \text{adom}(\mathcal{D})$, we from now on silently assume that L satisfies the latter two conditions.

The L -assignment $\text{as}_{\mathcal{I}}(L)$ realized in a model \mathcal{I} of \mathcal{O} with $L \subseteq \Delta^{\mathcal{I}}$ is the set that contains the following:

- $C(a)$ for all $C \in \text{cl}(\mathcal{O}, q)$ and $a \in L$ with $a \in C^{\mathcal{I}}$;
- $p(\bar{c})$ for all CQs $p(\bar{y}) \in \text{cl}(q)$ with $\bar{c} \in p(\mathcal{I})$ and $\bar{c} \in L$.

A k -subset of $\text{adom}(\mathcal{D})$ is a set $K \subseteq \text{adom}(\mathcal{D})$ such that $S \subseteq K$ and $|K \setminus S| \leq k$.

Let Γ be a set of L -assignments and let K be a k -subset. A Γ -choice for K is a subset $\Gamma' \subseteq \Gamma$ that contains exactly one L -assignment for each $L \subseteq K$. Every Γ -choice Γ' gives rise to a canonical database $\mathcal{D}_{\Gamma'}$ that is constructed by starting with $\mathcal{D}|_K$ and then doing the following for all $\mu \in \Gamma'$ and $p(\bar{c}) \in \mu$: rename all constants in $\mathcal{D}_p(\bar{c})$ except those in \bar{c} to fresh constants and add the resulting database. We say that Γ' is *consistent* if the following conditions are satisfied:

1. the extended database

$$\mathcal{D}|_K \cup \{C(a) \mid \exists \mu \in \Gamma' : C(a) \in \mu\}$$

is satisfiable with the \mathcal{ALCT}^u -ontology

$$\mathcal{O} \cup \{\top \sqsubseteq \neg p \mid p \text{ contraction of } q \text{ that is a BELIQ}\}.$$

2. If $\mathcal{D}_{\Gamma'} \models p(\bar{c})$ with $p(\bar{y}) \in \text{cl}(q)$ and $\bar{c} \in K$, and $\mu \in \Gamma'$ is an L -assignment with $\bar{c} \subseteq L$, then $p(\bar{c}) \in \mu$.

The algorithm is now exactly identical to the one used in the proof of Theorem 10. The running time can also be analyzed in a similar way, the only difference being that the number of L -assignments, for an $L \subseteq \text{adom}(\mathcal{D})$ with $|L| \leq \ell$, is now bounded single exponentially in $\|\mathcal{O}\|$ and *double* exponentially in $\|q\|$. The running time of the algorithm changes accordingly and yields fixed-parameter tractability with double exponentially running time. It remains to prove correctness.

Lemma 21. *The algorithm answers 'yes' iff $\mathcal{D}_{S,\ell,k}^{\approx} \models Q$.*

Proof. For the 'only if' direction, assume $\mathcal{D}_{S,\ell,k}^{\approx} \not\models Q$. We need to show that the algorithm answers 'no'. Since $\mathcal{D}_{S,\ell,k}^{\approx} \not\models Q$, there is a model \mathcal{I} of \mathcal{O} and $\mathcal{D}_{S,\ell,k}^{\approx}$ with $\mathcal{I} \not\models p$ for every disjunct p in q . For every $a \in \text{adom}(\mathcal{D}_{S,\ell,k}^{\approx})$, define

$$\text{tp}(a) = \{C \in \text{cl}(\mathcal{O}, q) \mid a \in C^{\mathcal{I}}\}$$

and for every $\bar{c} \subseteq \text{adom}(\mathcal{D}_{S,\ell,k}^{\approx})$, define

$$\text{inst}(\bar{c}) = \{p(\bar{y}) \in \text{cl}(q) \mid \bar{c} \in p(\mathcal{I})\}.$$

For every ℓ, k -sequence $v = S_0, O_0, \dots, O_{n-1}, S_n$ with $n \geq 1$ we define $L_v \subseteq \text{adom}(\mathcal{D})$ as $h(O_{n-1})$, where $h : \text{adom}(\mathcal{D}_{S,\ell,k}^{\approx}) \rightarrow \text{adom}(\mathcal{D})$ is the uncopying map, and we define the L_v -assignment

$$\mu_v = \{C(h(a)) \mid C \in a \in L_v, \text{tp}(a)\} \cup \{p(h(\bar{c})) \mid \bar{c} \subseteq L_v, p(\bar{y}) \in \text{inst}(\bar{c})\}.$$

Since \mathcal{I} is a model of $\mathcal{D}_{S,\ell,k}^{\approx}$ with $\mathcal{I} \not\models q$, μ_v indeed satisfies the conditions of L_v -assignments. Let $\Psi = \{\mu_v \mid v \text{ is an } \ell, k\text{-sequence of length } n \geq 1\}$. We argue by induction on i that $\Psi \subseteq \Gamma_i$ for every $i \geq 0$. Since Γ_0 is the set of all L -assignments, $\Psi \subseteq \Gamma_0$. For the induction step, assume $\Psi \subseteq \Gamma_i$ for some $i \geq 0$. Let $\mu_v \in \Psi$ be an L_v -assignment for some $v = S_0, O_0, \dots, O_{n-1}, S_n$. We show that for every k -subset K with $K \supseteq L_v$ there is a consistent Γ_i -choice Γ for K such that $\mu_v \in \Gamma$. Let K be a k -subset. For every $L \subseteq K$, choose a set $S_{n+1} \supseteq L$ with $|S_{n+1} \setminus S| \leq k$ and consider the sequence $v_{K,L} = S_0, O_0, \dots, O_{n-1}, K, L, S_{n+1}$. This is an ℓ, k -sequence that yields an L -assignment $\mu_{v_{K,L}} \in \Psi$. The set of all such L -assignments forms a Γ_i -choice Γ for K . Since \mathcal{I} is a model that realizes all L -assignments from Γ and such that $\mathcal{I} \not\models q$, Γ satisfies the two conditions of a consistent Γ_i -choice, so μ_v remains in Γ_{i+1} . Since Ψ is not empty, this implies that the algorithm answers 'no'.

For the 'if' direction, assume the algorithm answers 'no'. We need to show that $\mathcal{D}_{S,\ell,k}^{\approx} \not\models Q$, so we construct a model \mathcal{I} of \mathcal{O} and $\mathcal{D}_{S,\ell,k}^{\approx}$ with $\mathcal{I} \not\models q$. Since the algorithm answers 'no', the set $\Gamma := \Gamma_f$ is not empty. This implies that Γ contains at least one L -assignment for every set $L \subseteq \text{adom}(\mathcal{D})$ with $S \subseteq L$ and $|L \setminus S| \leq \ell$, since as soon as there is an L_0 such that some Γ_i does not contain an L_0 -assignment, the algorithm eventually eliminates all L -assignments and thus returns 'yes'. Thus, for every L -assignment $\mu \in \Gamma$ and every k -subset K , there is a consistent Γ -choice $\Gamma_{\mu,K}$ for K such that $\mu \in \Gamma_{\mu,K}$. We construct a model \mathcal{I} of $\mathcal{D}_{S,\ell,k}^{\approx}$ and \mathcal{O} with $\mathcal{I} \not\models q$.

For every ℓ, k -sequence $v = S_0, O_0, S_1, \dots, S_n$, we choose a consistent Γ -choice Γ_v , using the following inductive procedure. First, choose an arbitrary S -assignment $\mu_0 \in \Gamma$. For every sequence of the form $v = S_0$, choose a consistent Γ -choice $\Gamma_v \subseteq \Gamma$ with $\mu_0 \in \Gamma_v$. Since the algorithm did not remove μ_0 from Γ , there is such a Γ -choice Γ_v . For the induction step, assume that Γ_v has already been defined for all ℓ, k -sequences of length n and consider ℓ, k -sequence $v = S_0, O_0, \dots, S_n, O_n, S_{n+1}$ of length $n + 1$. Let $v' = S_0, O_0, \dots, S_n$. Since $O_n \subseteq S_n$, $\Gamma_{v'}$ contains a (unique) O_n -assignment μ . Choose a consistent Γ -choice $\Gamma_v \subseteq \Gamma$ for S_{n+1} . Such a Γ_v must exist, since the algorithm did not remove μ from Γ .

For every ℓ, k -sequence $v = S_0, O_0, \dots, S_n$, let \mathcal{D}_v be obtained from the canonical database \mathcal{D}_{Γ_v} by renaming every element $a \in \text{adom}(\mathcal{D}_{\Gamma_v}) \cap S_n$ to a_v , which is the name of the copy of a used in the bag for v in $\mathcal{D}_{S,\ell,k}^{\approx}$. For every element $a_v \in S_n$, let $\text{tp}(a) = \{C \in \text{cl}(\mathcal{O}, q) \mid C(a) \in \mu_v\}$. To construct \mathcal{I} , we start with the union of \mathcal{D}_v over all ℓ, k -sequences v . Additionally, for every $a_v \in \text{adom}(\mathcal{D}_{S,\ell,k})$, v an ℓ, k -sequence, choose a tree-shaped model \mathcal{I}_a of $\text{tp}(a)$ and add a copy of \mathcal{I}_a to \mathcal{I} and identify a with the root of \mathcal{I}_a . It can be verified that \mathcal{I} is indeed a model of \mathcal{O} . We refer to elements of $\Delta^{\mathcal{I}} \setminus \text{adom}(\mathcal{D}_{S,\ell,k})$ as *anonymous elements*.

To see that $\mathcal{I} \not\models q$, assume the opposite, so assume that there is a disjunct p of q and a homomorphism $h : p \rightarrow \mathcal{I}$. The range of h cannot only consist of anonymous elements, since p is connected and this would imply that the range of h is completely contained in one of the \mathcal{I}_a , which in turn

would imply $\text{tp}(a) \models p$, but contradicts the first condition of consistent Γ -choices. Thus, the range of h contains at least one element of $\text{adom}(\mathcal{D}_{S,\ell,k})$.

Since \mathcal{I} has treewidth (l, k) , there is a contraction p' of p of treewidth (l, k) such that there is an injective homomorphism $h' : p' \rightarrow \mathcal{I}$. For every ℓ, k -sequence $v = S_0, O_0, \dots, S_n$, let $B_v \subseteq \text{var}(p')$ be the set of all variables x such that $u(h'(x)) \in S_n$, where u is the uncopying map, or such that $h'(x)$ is an anonymous element in \mathcal{I}_a , where $u(a) \in S_n$. It is possible to choose a finite set V of ℓ, k -sequences such that $\bigcup_{v \in V} B_v = \text{var}(p')$ and such that V is connected, i.e.

- for each two $v, v' \in V$, V also contains the longest common prefix of v and v' and
- for each two $v, v' \in V$ with v' a prefix of v , V contains every ℓ, k -sequence v'' such that v' is a prefix of v'' and v'' is a prefix of v .

This yields an (ℓ, k) -tree decomposition (V, E, B_v) of p' , where $(v, v') \in E$ if and only if v' is an immediate successor of v . For every $v \in V$, let p'_v be p' restricted to variables that appear in B_v or in some $B_{v'}$ for some $v' \in V$ that is part of the subtree rooted at v . In particular, if v is the root of (V, E, B_v) , then $p'_v = p'$. For every p'_v where $v = S_0, O_0, \dots, S_{n-1}, O_{n-1}, S_n$ is not the root of (V, E, B_v) , make all variables x where $u(h'(x)) \in O_{n-1}$ answer variables and denote the resulting query by $p'_v(\bar{y}_v)$. Let $\bar{c}_v = u(h'(\bar{y}_v))$. If v is the root of (V, E, B_v) , we do not add any answer variables to p'_v .

Claim: For every $v \in V$, $p'_v(\bar{c}_v) \in \Gamma_v$.

The claim can be proven inductively on the position of v in (V, E, B_v) , starting with the leaves. If v is a leaf, then $p'_v(\bar{c}_v) \in \Gamma_v$ follows from the existence of the homomorphism h' restricted to $p'_v(\bar{y}_v)$. If $v = S_0, O_0, \dots, O_{n-1}, S_n$ is not a leaf, then let v_1, \dots, v_n be the children of v in (V, E, B_v) . By induction, the claim holds for v_1, \dots, v_n , so we have $p'_{v_i}(\bar{c}_{v_i}) \in \Gamma_{v_i}$ for all i with $1 \leq i \leq n$. Since every $c_{v_i} \subseteq S_n$, $p'_v(\bar{c}_v) \in \Gamma_v$ for all i with $1 \leq i \leq n$. Together with the homomorphism h' restricted to B_v , this yields $p'_v(\bar{c}_v) \in \Gamma_v$, which finishes the proof of the claim.

For v the root of V , the claim yields $p'() \in \Gamma_v$, which contradicts the requirement that every L -assignment $\mu \in \Gamma$ contains the negation of $p'()$ for every contraction p' of p . Thus, \mathcal{I} is a model of \mathcal{O} and $\mathcal{D}_{S,\ell,k}^{\approx}$ with $\mathcal{I} \not\models q$. \square

H Proofs for Section 8

H.1 Ontology strengthening approximation

Our aim is to prove Theorem 7. We split it here into two theorems, one for the upper bound and one for the lower bound. Let $\mathcal{ELIU}_{\perp}^{-}$ be the restriction of \mathcal{ELIU}_{\perp} where no disjunction is admitted on the left-hand side of concept inclusions. Note that every \mathcal{ELIU}_{\perp} -ontology can easily be transformed into an equivalent $\mathcal{ELIU}_{\perp}^{-}$ -ontology, but at the cost of an exponential blowup.

Theorem 11. *\mathcal{ELIU}_{\perp} -ontology strengthening OMQ evaluation in $(\mathcal{ELIU}_{\perp}, UCQ)$ is in 2EXPTIME in combined complexity and FPT with double exponential running time. In $(\mathcal{ELIU}_{\perp}^{-}, UCQ)$, it is in EXPTIME in combined complexity and FPT with single exponential running time.*

To prove Theorem 11, we start with some preliminaries. Let $\mathcal{I}_1, \dots, \mathcal{I}_n$ be interpretations. The direct product of $\mathcal{I}_1, \dots, \mathcal{I}_n$ is the interpretation \mathcal{I} with

$$\begin{aligned} \Delta^{\mathcal{I}} &= \Delta^{\mathcal{I}_1} \times \dots \times \Delta^{\mathcal{I}_n} \\ A^{\mathcal{I}} &= \{(a_1, \dots, a_n) \mid a_i \in A^{\mathcal{I}_i} \text{ for } 1 \leq i \leq n\} \\ r^{\mathcal{I}} &= \{((a_1, \dots, a_n), (b_1, \dots, b_2)) \mid (a_i, b_i) \in r^{\mathcal{I}_i} \\ &\quad \text{for } 1 \leq i \leq n\} \end{aligned}$$

for all role names r and concept names A . We denote the direct product of $\mathcal{I}_1, \dots, \mathcal{I}_n$ by $\prod_{i \in \{1, \dots, n\}} \mathcal{I}_i$, and we denote the direct product of two interpretations \mathcal{I}_1 and \mathcal{I}_2 by $\mathcal{I}_1 \otimes \mathcal{I}_2$. It is well known that being a model of an \mathcal{ELIU}_{\perp} ontology is preserved under direct products (Hernich et al. 2020), i.e. if $\mathcal{I}_1, \dots, \mathcal{I}_n$ are models of an \mathcal{ELIU}_{\perp} -ontology \mathcal{O} , then $\prod_{i \in \{1, \dots, n\}} \mathcal{I}_i$ is also a model of \mathcal{O} .

The proof of Theorem 11 is based on exhaustive \mathcal{ELIU}_{\perp} -approximation sets as introduced in the main body of the paper and crucially uses direct products. The following example gives a first idea.

Example 13. Let $\mathcal{O} = \{\top \sqsubseteq A_1 \sqcup A_2\} \cup \mathcal{O}'$, with \mathcal{O}' an \mathcal{ELIU} -ontology. Further let $\mathcal{O}_1 = \{\top \sqsubseteq A_1\} \cup \mathcal{O}'$ and $\mathcal{O}_2 = \{\top \sqsubseteq A_2\} \cup \mathcal{O}'$. We show below that $\mathcal{M} = \{\mathcal{O}_1, \mathcal{O}_2\}$ is an exhaustive \mathcal{ELIU}_{\perp} -approximation set. Consequently for all OMQs $Q(\bar{x}) = (\mathcal{O}, \Sigma, q)$ with q a UCQ, $\text{app}_{\mathcal{ELIU}_{\perp}}^{\uparrow}(Q, \mathcal{D}) = Q_1(\mathcal{D}) \cap Q_2(\mathcal{D})$ where $Q_i = (\mathcal{O}_i, \Sigma, q)$.

Assume to the contrary that \mathcal{M} is not an exhaustive \mathcal{ELIU}_{\perp} -approximation set. Then there is an \mathcal{ELIU}_{\perp} -ontology \mathcal{O}' such that $\mathcal{O}' \models \mathcal{O}$, $\mathcal{O}' \not\models \mathcal{O}_1$, and $\mathcal{O}' \not\models \mathcal{O}_2$. This means that there are models \mathcal{I}_1 and \mathcal{I}_2 of \mathcal{O}' , such that \mathcal{I}_i is not a model of \mathcal{O}_i for $i \in \{1, 2\}$. In particular, there is an element $a \in \Delta^{\mathcal{I}_1}$ such that $a \notin A^{\mathcal{I}_1}$. Then $\mathcal{I}_1 \otimes \mathcal{I}_2$ is a model of \mathcal{O}' , but $(a, b) \in (\neg A_1 \sqcap \neg A_2)^{\mathcal{I}_1 \otimes \mathcal{I}_2}$. Thus, $\mathcal{I}_1 \otimes \mathcal{I}_2$ is not a model of \mathcal{O} , contradicting $\mathcal{O}' \models \mathcal{O}$.

For a given \mathcal{ELIU}_{\perp} -concept C , define the set $f(C)$ inductively as follows:

- If $C \in \{\top, \perp\} \cup \text{NC}$, then $f(C) = \{C\}$.
- If $C = \exists r.D$, then $f(C) = \{\exists r.D' \mid D' \in f(D)\}$.
- If $C = D_1 \sqcap D_2$, then $f(C) = \{D'_1 \sqcap D'_2 \mid D'_1 \in f(D_1), D'_2 \in f(D_2)\}$.
- If $C = D_1 \sqcup D_2$, then $f(C) = f(D_1) \cup f(D_2)$.

It is easy to check that every \mathcal{ELIU}_{\perp} -concept C is equivalent to $\bigsqcup_{C' \in f(C)} C'$.

Given an \mathcal{ELIU}_{\perp} -ontology \mathcal{O} , we first construct an equivalent ontology \mathcal{O}' obtained from \mathcal{O} by replacing every $C \sqsubseteq D \in \mathcal{O}$ by the set of CIs $\{C' \sqsubseteq D \mid C' \in f(C)\}$. If \mathcal{O} is formulated in $\mathcal{ELIU}_{\perp}^{-}$, then $\mathcal{O}' = \mathcal{O}$. In general though, the size of \mathcal{O}' is exponential in that of \mathcal{O} . Now we define $\mathcal{E}_{\mathcal{O}}$ to be the set of ontologies \mathcal{O}'' that can be obtained by choosing for every CI $C \sqsubseteq D \in \mathcal{O}'$, a CI $C \sqsubseteq D'$ with $D' \in f(D)$ and including it in \mathcal{O}'' .

Lemma 22. For every \mathcal{ELIU}_\perp -ontology \mathcal{O} , the set $\mathcal{E}_\mathcal{O}$ is an exhaustive \mathcal{ELI}_\perp -approximation set for \mathcal{O} .

Proof. Assume $\mathcal{E}_\mathcal{O}$ is not an exhaustive \mathcal{ELI}_\perp -approximation set for \mathcal{O} . Then there exists an \mathcal{ELI}_\perp -ontology \mathcal{O}' such that $\mathcal{O}' \models \mathcal{O}$ and $\mathcal{O}' \not\models \widehat{\mathcal{O}}$ for every $\widehat{\mathcal{O}} \in \mathcal{E}_\mathcal{O}$. Thus, for every $\widehat{\mathcal{O}} \in \mathcal{E}_\mathcal{O}$ there is a model $\mathcal{I}_{\widehat{\mathcal{O}}}$ of \mathcal{O}' that is not a model of $\widehat{\mathcal{O}}$. Thus, for every $\widehat{\mathcal{O}}$ there is a CI $C_{\widehat{\mathcal{O}}} \sqsubseteq D_{\widehat{\mathcal{O}}} \in \widehat{\mathcal{O}}$ that is violated by $\mathcal{I}_{\widehat{\mathcal{O}}}$. Let S be the set of all such CIs that are violated by some $\mathcal{I}_{\widehat{\mathcal{O}}}$. We claim that there is a $C \sqsubseteq D \in \mathcal{O}$ such that S contains $C \sqsubseteq D'$ for all $D' \in f(D)$. If this was not the case, then for every $C \sqsubseteq D \in \mathcal{O}$, there is a $D' \in f(D)$ with $C \sqsubseteq D' \notin S$. But then the ontology $\mathcal{O}_0 \in \mathcal{E}_\mathcal{O}$ obtained as the product of these choices contains only CIs that are never violated in any of the $\mathcal{I}_\mathcal{O}$, contradicting that $\mathcal{I}_{\mathcal{O}_0}$ violates some CI from \mathcal{O}_0 . From now on, let $C \sqsubseteq D \in \mathcal{O}$ a fixed CI that fulfils $\{C \sqsubseteq D' \mid D' \in f(D)\} \subseteq S$. Let $\mathcal{M} = \{\widehat{\mathcal{O}} \in \mathcal{E}_\mathcal{O} \mid \mathcal{I}_{\widehat{\mathcal{O}}}$ violates a CI of the form $C \sqsubseteq D'$ for some $D' \in f(D)\}$.

Define $\mathcal{I} = \prod_{\widehat{\mathcal{O}} \in \mathcal{M}} \mathcal{I}_{\widehat{\mathcal{O}}}$. Since models of \mathcal{ELI}_\perp -ontologies are closed under products, \mathcal{I} is a model of \mathcal{O}' . We aim to show that \mathcal{I} is not a model of \mathcal{O} , contradicting $\mathcal{O}' \models \mathcal{O}$. For every $\widehat{\mathcal{O}} \in \mathcal{M}$, $\mathcal{I}_{\widehat{\mathcal{O}}}$ violates a CI of the form $C \sqsubseteq D'$ for some $D' \in f(D)$. Let $a_{\widehat{\mathcal{O}}} \in \Delta^{\mathcal{I}_{\widehat{\mathcal{O}}}}$ such that $a_{\widehat{\mathcal{O}}} \in \mathcal{I}^{C_{\widehat{\mathcal{O}}}} \setminus \mathcal{I}^{D'_{\widehat{\mathcal{O}}}}$. Since models of \mathcal{ELI}_\perp -ontologies are closed under products, the element $a_0 := \prod_{\widehat{\mathcal{O}} \in \mathcal{M}} a_{\widehat{\mathcal{O}}}$ fulfils $a_0 \in \mathcal{I}^C$, but $a_0 \notin \mathcal{I}^{D'}$ for all $D' \in f(D)$. Since $D \equiv \bigsqcup_{D' \in f(D)} D'$, $a_0 \notin \mathcal{I}^D$, showing that \mathcal{I} is not a model of \mathcal{O} . \square

Lemma 22 allows us to compute $\text{app}_{\mathcal{ELI}_\perp}^\uparrow(Q, \mathcal{D})$ for any given $Q \in (\mathcal{ELIU}_\perp, \text{UCQ})$ and database \mathcal{D} by computing $\mathcal{E}_\mathcal{O}$ and evaluating the query under every ontology $\widehat{\mathcal{O}} \in \mathcal{E}_\mathcal{O}$. More precisely, we use the following algorithm: Given $Q = (\mathcal{O}, \Sigma, q(\bar{x})) \in (\mathcal{ELIU}_\perp, \text{UCQ})$, a Σ -database \mathcal{D} , and $\bar{a} \in \text{dom}(\mathcal{D})^{|\bar{x}|}$, iterate over all ontologies $\mathcal{O}' \in \mathcal{E}_\mathcal{O}$ and check whether $\mathcal{O}', \mathcal{D} \models Q(\bar{a})$. If this is the case for every $\mathcal{O}' \in \mathcal{E}_\mathcal{O}$, answer *yes*, otherwise *no*. Correctness follows from Lemma 22 and the definition of \mathcal{ELI}_\perp -approximation sets. For the running time, note that the number of ontologies in $\mathcal{E}_\mathcal{O}$ is at most double exponential in $\|\mathcal{O}\|$ and every $\mathcal{O}' \in \mathcal{E}_\mathcal{O}$ is of size at most single exponential in $\|\mathcal{O}\|$. Since answering OMQs from $(\mathcal{ELI}_\perp, \text{UCQ})$ is in EXPTIME (Eiter et al. 2008) and FPT with single exponential running time, this yields a 2EXPTIME algorithm and FPT with double exponential running time. If the original ontology \mathcal{O} is formulated in \mathcal{ELIU}_\perp , then the ontologies in $\mathcal{E}_\mathcal{O}$ are only of polynomial size and thus we obtain an EXPTIME algorithm.

Theorem 12. \mathcal{ELI}_\perp -ontology strengthening OMQ evaluation in $(\mathcal{ELIU}_\perp, \text{AQ})$ is 2EXPTIME-hard in combined complexity.

We prove Theorem 12 by a reduction from the word problem of alternating exponentially space bounded Turing machines, which is 2EXPTIME-complete (Chandra, Kozen, and Stockmeyer 1981). An *alternating Turing Machine (ATM)* is a tuple $(Q, \Sigma, \Gamma, q_0, q_a, q_r, \Delta)$ where

- $Q = Q_\exists \uplus Q_\forall \uplus \{q_a, q_r\}$ is the set of states,
- Σ is the input alphabet,
- $\Gamma \supseteq \Sigma$ is the tape alphabet that contains a *blank symbol* $\square \in \Gamma \setminus \Sigma$,
- $q_0 \in Q$ the starting state,
- $q_a \in Q$ the accepting state,
- $q_r \in Q$ the rejecting state and
- $\Delta \subseteq (Q_\exists \cup Q_\forall) \times \Gamma \times \Gamma \times \{l, r\} \times Q$ the transition relation.

A *configuration* is a word wqw' with $w, w' \in \Gamma^*$ and $q \in Q$. The intended meaning is that the tape contains the word ww' (with only blanks before and behind it), the machine is in state q and the head is on the leftmost symbol of w' . A configuration vqv' is a *successor configuration* of a configuration $wq\gamma_1w'$ if there is a tuple $(q, \gamma_1, \gamma_2, d, p) \in \Delta$ and vqv' is obtained from $wq\gamma_1w'$ by replacing the γ_1 under the head by γ_2 , change the state from q to p and move the head one step into direction d . A *halting configuration* is of the form vqv' with $q \in \{q_a, q_r\}$. We inductively define accepting configurations. Let c be a configuration.

- If c is of the form $wq_a w'$, then c is *accepting*.
- If c is of the form wqw' with $q \in Q_\exists$ and there is a successor configuration c' of c such that c' is accepting, then c is *accepting*.
- If c is of the form wqw' with $q \in Q_\forall$ and every successor configuration c' of c is accepting, then c is *accepting*.

The *starting configuration* for a word $w \in \Sigma^*$ is the configuration q_0w . A word w is accepted by M if the starting configuration for w is accepting. We use $L(M)$ to denote the set $\{w \in \Sigma^* \mid M \text{ accepts } w\}$.

Let M be an exponentially space bounded ATM that decides a 2EXPTIME-complete problem. Given an input word $w \in \Sigma^*$, we construct in polynomial time an OMQ $Q \in (\mathcal{ELIU}_\perp, \text{AQ})$ and a database $\mathcal{D} = \{A(a)\}$ such that $w \in L(M)$ if and only if $a \in \text{app}_{\mathcal{ELI}_\perp}^\uparrow(Q, \mathcal{D})$. W.l.o.g. we make the following assumptions about M :

- $q_0 \in Q_\exists$.
- For every $q \in Q_\exists \cup Q_\forall$ and $\gamma_1 \in \Gamma$, there are precisely two transitions of the form $(q, \gamma_1, \gamma_2, d, p) \in \Delta$. We consider Δ as two functions $\delta_1, \delta_2 : (Q_\exists \cup Q_\forall) \times \Gamma \rightarrow \Gamma \times \{l, r\} \times Q$ where each δ_i yields one of the two transitions.
- For every configuration with a state from Q_\exists , both successor configurations have a state from Q_\forall , and vice versa.
- On an input w , M uses at most $2^{|w|}$ cells of the tape, and it does not use cells to the left of the starting position of the head.

The general idea for the reduction is to construct \mathcal{O} in such a way that of all the ontologies in $\mathcal{E}_\mathcal{O}$, only a single ontology \mathcal{O}_0 is satisfiable with \mathcal{D} . By Lemma 22, this implies that $a \in \text{app}_{\mathcal{ELI}_\perp}^\uparrow(Q, \mathcal{D})$ if and only if $\mathcal{O}_0, \mathcal{D} \models Q(a)$. The ontology \mathcal{O}_0 has the effect that starting from the single assertion $A(a)$ in the database, it creates the tree of all relevant configurations of M on input w in the anonymous part,

and if M accepts w , it propagates the query concept name B back to the element a .

We describe the reduction in two steps. First, we describe \mathcal{O}_0 and argue that $\mathcal{O}_0, \mathcal{D} \models Q(a)$ if and only if $w \in L(M)$. Secondly, we construct \mathcal{O} such that $\mathcal{O}_0 \in \mathcal{E}_{\mathcal{O}}$ and such that all other ontologies in $\mathcal{E}_{\mathcal{O}}$ are unsatisfiable with \mathcal{D} .

Let $w = w_1 w_2 \dots w_n \in \Sigma^*$. Starting from A , we generate an infinite binary tree where each node stands for a configuration. The concept name T represents a configuration, the concept name T_{\exists} represents configurations where the state is from Q_{\exists} , and the concept name T_{\forall} represents configurations where the state is from Q_{\forall} . The role name s serves as the successor relation, and we use concept names S_1 and S_2 to distinguish the two successors from each other.

$$\begin{aligned} A &\sqsubseteq T \sqcap T_{\exists} \\ T &\sqsubseteq \exists s.(T \sqcap S_1) \\ T &\sqsubseteq \exists s.(T \sqcap S_2) \\ \exists s^-.T_{\exists} &\sqsubseteq T_{\forall} \\ \exists s^-.T_{\forall} &\sqsubseteq T_{\exists} \end{aligned}$$

To store a configuration of M in every node, we let T generate a binary tree of depth n , so that the 2^n leaves of this tree can be used as tape cells. We use concept names of the form B_i^j , $0 \leq i < n$, $j \in \{0, 1\}$, to encode which leaf encodes which position of the tape, where B_i^j means that the i th bit of the position (in binary) is j . For convenience, we label the two children of every inner node of the tree with L and R to indicate whether the subtree contains the left or the right half of the tape.

$$\begin{aligned} T &\sqsubseteq L_0 \\ L_i &\sqsubseteq \exists r.(L_{i+1} \sqcap B_i^0 \sqcap L) \quad \text{for } 0 \leq i < n \\ L_i &\sqsubseteq \exists r.(L_{i+1} \sqcap B_i^1 \sqcap R) \quad \text{for } 0 \leq i < n \\ \exists r^-.B_i^j &\sqsubseteq B_i^j \quad \text{for } 0 \leq i < n, j \in \{0, 1\} \end{aligned}$$

Next, we populate the tree below A with the starting configuration. We first label every node in the tree of the starting configuration with a concept name T_0 .

$$\begin{aligned} A &\sqsubseteq T_0 \\ \exists r^-.T_0 &\sqsubseteq T_0 \end{aligned}$$

The content of a tape cell is encoded in a concept name of the form $M_{\gamma,q}$ or M_{γ} , where $\gamma \in \Gamma$ and $q \in Q$. A concept name of the form $M_{\gamma,q}$ means that the tape cell contains the symbol γ the head is currently in this tape cell, and the ATM is in state q . A concept name M_{γ} means that the tape cell contains the symbol γ and the head is currently not in this tape cell. For every $x \in \{0, \dots, 2^n - 1\}$ with binary representation $x_0 x_1 \dots x_{n-1}$, we write B_x for the concept $B_0^{x_0} \sqcap B_1^{x_1} \sqcap \dots \sqcap B_{n-1}^{x_{n-1}}$. We introduce the following CIs:

$$\begin{aligned} T_0 \sqcap L_n \sqcap B_0 &\sqsubseteq M_{w_0, q_0} \\ T_0 \sqcap L_n \sqcap B_i &\sqsubseteq M_{w_i} \quad \text{for } 1 \leq i \leq n \end{aligned}$$

The next CIs write blanks into all of the remaining positions.

$$\begin{aligned} T_0 \sqcap L_n \sqcap B_n &\sqsubseteq E_0 \\ \exists r.E_0 &\sqsubseteq E_0 \\ R \sqcap \exists r^-. \exists r.(L \sqcap E_0) &\sqsubseteq E_1 \\ \exists r^-.E_1 &\sqsubseteq E_1 \\ E_1 \sqcap L_n &\sqsubseteq M_{\square} \end{aligned}$$

This finishes the starting configuration. Next, we implement the transitions of M . Let

$$\mathcal{M} = \{M_{\gamma,q}, M_{\gamma} \mid \gamma \in \Gamma, q \in Q\}.$$

For every concept name $M \in \mathcal{M}$, we introduce a concept name M' . To implement the transitions of the ATM, we first copy every configuration to its two successor configurations, but using the primed concept names in the successor. For every number $i \in \{0, \dots, 2^n - 1\}$ and every $M \in \mathcal{M}$, use the following CIs:

$$B_i \sqcap \exists (r^-)^n. \exists s^-. \exists r^n.(B_i \sqcap M) \sqsubseteq M' \quad (1)$$

A short remark regarding the second part of the proof: Note that the number of CIs introduced in (1) is exponential in n . In the second part of the proof, we will show how to produce these CIs in \mathcal{O}_0 using only a single CI in \mathcal{O} that can be replaced in exponentially many ways. The key problem of the second part of the proof will then be to construct \mathcal{O} such that all other ontologies in $\mathcal{E}_{\mathcal{O}}$ that are obtained by replacing the CI in \mathcal{O} in an unintended way become unsatisfiable with \mathcal{D} .

Now we calculate the two successor configurations that are determined by δ_1 and δ_2 . We propagate the concept names S_1 and S_2 to the leaves of the trees by introducing for $i \in \{1, 2\}$ the following CI:

$$\exists r^-.S_i \sqsubseteq S_i$$

When computing the i th successor configuration for $i \in \{1, 2\}$, the content of the cell that contained the head in the previous configuration depends on δ_i . So for every $\gamma \in \Gamma$ and $q \in Q$ and $i \in \{1, 2\}$, let $\delta_i(q, \gamma) = (\alpha, d, p)$ for some $\alpha \in \Gamma$, $d \in \{l, r\}$ and $p \in Q$. We introduce the following CI:

$$M'_{\gamma,q} \sqcap S_i \sqsubseteq M_{\alpha} \sqcap F_{p,d}$$

Here, $F_{p,d}$ is a fresh concept name that indicates that a transition needs to be done, namely into state p while moving the head one step in direction d . We propagate $F_{p,d}$ to the leaf that represents the tape cell one step in direction d . For the case $d = l$, we introduce the following CIs for every $\beta \in \Gamma$:

$$\begin{aligned} \exists r.(L \sqcap F_{p,l}) &\sqsubseteq F_{p,l} \\ L \sqcap \exists r^-. \exists r.(R \sqcap F_{p,l}) &\sqsubseteq F'_{p,l} \\ R \sqcap \exists r^-. F'_{p,l} &\sqsubseteq F'_{p,l} \\ F'_{p,l} \sqcap M'_{\beta} \sqcap L_n &\sqsubseteq M_{\beta,p} \end{aligned}$$

For the case $d = r$, we introduce the following CIs for every $\beta \in \Gamma$:

$$\begin{aligned} \exists r.(R \sqcap F_{p,r}) &\sqsubseteq F_{p,r} \\ R \sqcap \exists r^-. \exists r.(L \sqcap F_{p,r}) &\sqsubseteq F'_{p,r} \\ L \sqcap \exists r^-. F'_{p,r} &\sqsubseteq F'_{p,r} \\ F'_{p,r} \sqcap M'_{\beta} \sqcap L_n &\sqsubseteq M_{\beta,p} \end{aligned}$$

Finally, we send a marker H to every other tape cell, to notify the tape cell that it does not contain the head of the ATM in this configuration. For every $\gamma \in \Gamma$ and every $q \in Q$, we introduce the CI

$$M_{\gamma,q} \sqsubseteq H$$

and propagate it to all other tape cells using the following CIs:

$$\begin{aligned} \exists r.H &\sqsubseteq H \\ R \sqcap \exists r^-. \exists r.(L \sqcap H) &\sqsubseteq H' \\ L \sqcap \exists r^-. \exists r.(R \sqcap H) &\sqsubseteq H' \\ \exists r^-. H' &\sqsubseteq H' \end{aligned}$$

After the marker arrived, we can create the symbol that encodes the content of the tape cell. For every $\beta \in \Gamma$, we introduce the following CIs:

$$H' \sqcap M'_\beta \sqcap L_n \sqsubseteq M_\beta$$

To check acceptance, we mark the root of every tree that encodes an accepting configuration with the concept name B . For every $\gamma \in \Gamma$, introduce the following CIs:

$$\begin{aligned} M_{\gamma,q_a} &\sqsubseteq B \\ \exists r.B &\sqsubseteq B \\ T_\forall \sqcap \exists s.(S_1 \sqcap B) \sqcap \exists s.(S_2 \sqcap B) &\sqsubseteq B \\ T_\exists \sqcap \exists s.B &\sqsubseteq B \end{aligned}$$

This finishes the definition of \mathcal{O}_0 . It can be verified that $w \in L(M)$ if and only if $\mathcal{O}_0, \mathcal{D} \models B(a)$.

For the second part of the proof, we would like to construct the ontology \mathcal{O} such that \mathcal{O} is of polynomial size, $\mathcal{O}_0 \in \mathcal{E}_{\mathcal{O}}$ and such that \mathcal{D} is unsatisfiable with every $\mathcal{O}' \in \mathcal{E}_{\mathcal{O}}$, $\mathcal{O}' \neq \mathcal{O}_0$. In fact, our construction of \mathcal{O} will not achieve this, but instead produces an ontology $\mathcal{O}_1 \in \mathcal{E}_{\mathcal{O}}$ which satisfies $\mathcal{O}_1 \supseteq \mathcal{O}_0$, and such that \mathcal{O}_0 and \mathcal{O}_1 yield the same certain answers to OMQs on \mathcal{D} , which is sufficient to make the reduction work. The ontology \mathcal{O} consists of all CIs in \mathcal{O}_0 , besides the CIs (1). Instead of the exponentially many CIs (1), we include in \mathcal{O} the following single CI.

$$\begin{aligned} &(B_0^0 \sqcup B_0^1) \sqcap \dots \sqcap (B_{n-1}^0 \sqcup B_{n-1}^1) \sqcap \\ &(\exists r^-.)^n \exists s^-. (\exists r.)^n (\\ &(B_0^0 \sqcup B_0^1) \sqcap \dots \sqcap (B_{n-1}^0 \sqcup B_{n-1}^1) \sqcap \bigsqcup_{M \in \mathcal{M}} M) \\ &\sqsubseteq (\bigsqcup_{M \in \mathcal{M}} M') \sqcup D \end{aligned} \quad (2)$$

Here, D is a fresh dummy concept name. Recall the procedure how $\mathcal{E}_{\mathcal{O}}$ is obtained from an ontology \mathcal{O} and note that applying the construction of $\mathcal{E}_{\mathcal{O}}$ to the CI (2) indeed yields an ontology that contains (among others) all the CIs (1). In fact, the construction of $\mathcal{E}_{\mathcal{O}}$ forces us to include a CI for every concept from $f(\text{LHS of (2)})$. This also includes concepts where the two encoded tape positions in the left hand side are different, i.e. where there exists an $i \in \{0, \dots, n-1\}$ such that once B_i^0 was chosen and once B_i^1 was chosen to appear in the CI. For all CIs of this form, we would like to choose the dummy concept name D on the right hand side.

To summarize, there are two kinds unintended CIs that we want to ‘disable’, meaning that we want every ontology from $\mathcal{E}_{\mathcal{O}}$ containing such a CI to become unsatisfiable with \mathcal{D} . The first kind are CIs of the form

$$B_i \sqcap (\exists r^-.)^n \exists s^-. (\exists r.)^n (B_i \sqcap M) \sqsubseteq N'$$

with $M, N \in \mathcal{M}$ and $M \neq N$. The second kind are of the form

$$B_i \sqcap (\exists r^-.)^n \exists s^-. (\exists r.)^n (B_j \sqcap M) \sqsubseteq N'$$

with $M, N \in \mathcal{M}$ and $i \neq j$. To achieve this, for every unintended CI, we generate a substructure that contradicts this CI.

For every concept from $f(\text{LHS of (2)})$, we generate an element that satisfies that concept. We construct a tree using a fresh role name t such that for every $i \in \{0, \dots, n-1\}$, every leaf is labelled with precisely one concept name from X_i^0 and X_i^1 , precisely one concept name from Y_i^0 and Y_i^1 , and precisely one concept name of the form \widehat{M} for some $M \in \mathcal{M}$. We use concept names V_i , $i \in \{0, \dots, 2n+1\}$ to label the layers of the tree. For all $i \in \{0, \dots, n-1\}$ and $j \in \{0, 1\}$, introduce the following CIs:

$$\begin{aligned} A &\sqsubseteq V_0 \\ V_i &\sqsubseteq \exists t.(V_{i+1} \sqcap X_i^0) \sqcap \exists t.(V_{i+1} \sqcap X_i^1) \\ V_{n+i} &\sqsubseteq \exists t.(V_{n+i+1} \sqcap Y_{n+i}^0) \sqcap \exists t.(V_{n+i+1} \sqcap Y_{n+i}^1) \\ \exists t^-. X_i^j &\sqsubseteq X_i^j \\ \exists t^-. Y_i^j &\sqsubseteq Y_i^j \end{aligned}$$

The last layer of the tree is generated by the following CIs for every $M \in \mathcal{M}$.

$$V_{2n} \sqsubseteq \exists t.(V_{2n+1} \sqcap \widehat{M})$$

At every leaf, we figure out whether the two numbers encoded by the X_i^j and the Y_i^j are equal. If so, we derive a marker $K_=$, otherwise a marker K_\neq . For every $i \in \{0, \dots, n-1\}$, introduce the following CIs:

$$\begin{aligned} X_i^0 \sqcap Y_i^1 &\sqsubseteq K_\neq \\ X_i^1 \sqcap Y_i^0 &\sqsubseteq K_\neq \\ X_i^0 \sqcap Y_i^0 &\sqsubseteq Z_i \\ Z_0 \sqcap \dots \sqcap Z_{n-1} &\sqsubseteq K_= \end{aligned}$$

Below every leaf of the tree, we generate a substructure used to contradict one specific CI. For every $i \in \{0, \dots, n-1\}$ and $j \in \{0, 1\}$, introduce the following CIs:

$$\begin{aligned} V_{2n+1} &\sqsubseteq \exists (r^-)^n. \exists s^-. \exists r^n. V \\ \exists (r^-)^n. \exists s. \exists r^n. X_i^j &\sqsubseteq X_i^j \\ \exists (r^-)^n. \exists s. \exists r^n. \widehat{M} &\sqsubseteq \widehat{M} \\ Y_i^j \sqcap V_{2n+1} &\sqsubseteq B_i^j \\ X_i^j \sqcap V &\sqsubseteq B_i^j \\ \widehat{M} \sqcap V &\sqsubseteq M \end{aligned}$$

Here, the element satisfying V_{2n+1} plays the role of a tape cell in a successor configuration, while the element satisfying V plays the role of a tape cell in the previous configuration. To disable all CIs of the form

$$B_i \sqcap \exists(r^-)^n . \exists s^- . \exists r^n . (B_i \sqcap M) \sqsubseteq N'$$

with $M \neq N$, we introduce for every $M \in \mathcal{M}$ the CI

$$K_{=} \sqcap \widehat{M} \sqsubseteq M'$$

and for each two $M_1, M_2 \in \mathcal{M}$ with $M_1 \neq M_2$, we introduce the CIs

$$\begin{aligned} K_{=} \sqcap M'_1 \sqcap M'_2 &\sqsubseteq \perp \\ K_{=} \sqcap M'_1 \sqcap D &\sqsubseteq \perp \end{aligned}$$

To disable all CIs of the form

$$B_i \sqcap \exists(r^-)^n . \exists s^- . \exists r^n . (B_j \sqcap M) \sqsubseteq N'$$

with $i \neq j$, we have introduced the dummy concept name D in the right hand side of CI (2). We introduce for every $M \in \mathcal{M}$ the following CI:

$$K_{\neq} \sqcap M' \sqsubseteq \perp$$

This finishes the definition of \mathcal{O} . It can be verified that in $\mathcal{E}_{\mathcal{O}}$, the only ontology that is satisfiable with \mathcal{D} is the unique ontology \mathcal{O}_1 that contains every CI from \mathcal{O}_0 . This yields the following result.

Lemma 23. *For the OMQ $Q = (\mathcal{O}, \{A\}, B(x))$ and database $\mathcal{D} = \{A(a)\}$ constructed from M and w , it holds that $w \in L(M)$ if and only if $a \in \text{app}_{\mathcal{E}\mathcal{L}\mathcal{I}_{\perp}}^{\uparrow}(Q, \mathcal{D})$.*

We discuss Example 11 in which $\mathcal{O} = \{\exists r . \top \sqcap \forall r . A \sqsubseteq B_1 \sqcup B_2\}$. Then for each $n \geq 1$, the $\mathcal{E}\mathcal{L}\mathcal{I}_{\perp}$ -ontology

$$\mathcal{O}_n = \{\exists r . A \sqsubseteq \exists r^n . X, \exists r . (A \sqcap \exists r^{n-1} . X) \sqsubseteq B_1\}$$

is such that $\mathcal{O}_n \models \mathcal{O}$. We show that every $\mathcal{E}\mathcal{L}\mathcal{I}_{\perp}$ -ontology \mathcal{O}'_n with $\mathcal{O}_n \models \mathcal{O}'_n \models \mathcal{O}$ is equivalent to \mathcal{O}_n . Assume \mathcal{O}'_n with $\mathcal{O}_n \models \mathcal{O}'_n \models \mathcal{O}$ is given and assume for a proof by contradiction that $\mathcal{O}'_n \not\models \mathcal{O}_n$. Assume first that $\mathcal{O}'_n \not\models \exists r . A \sqsubseteq \exists r^n . X$. Then one can show that the interpretation \mathcal{I} with $\Delta^{\mathcal{I}} = \{0, \dots, n\}$ and $r^{\mathcal{I}} = \{(i, i+1) \mid 0 \leq i < n\}$, $A^{\mathcal{I}} = \{1\}$ is a model of \mathcal{O}'_n and we have derived a contradiction to $\mathcal{O}'_n \models \mathcal{O}$. Now assume that $\mathcal{O}'_n \not\models \exists r . (A \sqcap \exists r^{n-1} . X) \sqsubseteq B_1$. Then one can show that the interpretation \mathcal{I}' obtained from \mathcal{I} by adding n to $X^{\mathcal{I}}$ is a model of \mathcal{O}'_n and we have again derived a contradiction to $\mathcal{O}'_n \models \mathcal{O}$.

The following example shows that transforming an $\mathcal{ALCC}\mathcal{I}$ -ontology into an $\mathcal{E}\mathcal{L}\mathcal{I}\mathcal{U}_{\perp}$ -ontology in a straightforward way does not preserve $\mathcal{E}\mathcal{L}\mathcal{I}_{\perp}$ -ontology strengthening OMQ evaluation.

Example 14. *Consider the ontology*

$$\begin{aligned} \mathcal{O} = \{ &\top \sqsubseteq \exists r . \neg A \sqcup \exists s . \neg B, \\ &\top \sqsubseteq \exists r . \top \sqcap \exists s . \top, \\ &\exists r^- . E \sqsubseteq A, \exists s^- . F \sqsubseteq B \} \end{aligned}$$

and the database $\mathcal{D} = \{E(a), F(b)\}$. A typical procedure to obtain an $\mathcal{E}\mathcal{L}\mathcal{I}\mathcal{U}_{\perp}$ -ontology \mathcal{O}' from \mathcal{O} introduces fresh

concept names A' and B' , adds the CIs $A' \sqsubseteq \neg A$, and $B' \sqsubseteq \neg B$ to \mathcal{O} , and then replaces $\neg A$ by A' and $\neg B$ by B' in \mathcal{O} . As \mathcal{D} is not satisfiable w.r.t. the ontologies obtained from \mathcal{O}' by replacing $\top \sqsubseteq \exists r . A' \sqcup \exists s . B'$ by $\top \sqsubseteq \exists r . A'$ or by $\top \sqsubseteq \exists s . B'$, we obtain that $a \in \text{app}_{\mathcal{E}\mathcal{L}\mathcal{I}_{\perp}}^{\uparrow}(Q, \mathcal{D})$ for any $Q(x) = (\mathcal{O}', \Sigma, q(x))$.

This is not the case for \mathcal{O} itself. For the ontology \mathcal{O}^* obtained from \mathcal{O} by adding the CIs $A' \sqsubseteq \neg A$, and $B' \sqsubseteq \neg B$ and replacing $\top \sqsubseteq \exists r . \neg A \sqcup \exists s . \neg B$ by $\exists r . (M \sqcap A) \sqsubseteq \exists s . B'$ and $\top \sqsubseteq \exists r . M$ we have that $\mathcal{O}^* \models \mathcal{O}$ but \mathcal{D} is satisfiable w.r.t. \mathcal{O}^* .

H.2 Database strengthening approximation

For classes of pointed databases \mathfrak{D} that are closed under substructures, \mathfrak{D} -database strengthening approximations can be characterized in a natural way. For $\langle \mathcal{D}, \bar{a} \rangle \in \mathfrak{D}$, let $\text{id}_{\mathfrak{D}}(\mathcal{D}, \bar{a})$ be the set of pointed databases $\langle \mathcal{D}', \bar{a}' \rangle$ that can be obtained from $\langle \mathcal{D}, \bar{a} \rangle \in \mathfrak{D}$ by identifying constants. Every such identification gives rise to an obvious *provenance homomorphism* h from \mathcal{D} to \mathcal{D}' with $h(\bar{a}) = \bar{a}'$.

Lemma 24. *Let \mathfrak{D} be a class of pointed databases that is closed under substructures, $Q(\bar{x}) = (\mathcal{O}, \Sigma, q) \in (\text{FO}, \text{UCQ})$ an OMQ, \mathcal{D} a Σ -database and $\bar{a} \in \text{adom}(\mathcal{D})^{|\bar{x}|}$. Then $\bar{a} \in \text{app}_{\mathfrak{D}}^{\uparrow}(Q, \mathcal{D})$ iff for all $\langle \mathcal{D}', \bar{a}' \rangle \in \text{id}_{\mathfrak{D}}(\mathcal{D}, \bar{a})$, $\bar{a}' \in Q(\mathcal{D}')$.*

Proof. Only the ‘if’ needs to be proved. Assume that for all $\langle \mathcal{D}', \bar{a}' \rangle \in \text{id}_{\mathfrak{D}}(\mathcal{D}, \bar{a})$, $\bar{a}' \in Q(\mathcal{D}')$. Let $\langle \mathcal{D}', \bar{b} \rangle \in \mathfrak{D}$ and h be a homomorphism from \mathcal{D} to \mathcal{D}' with $h(\bar{a}) = \bar{b}$. Let $\mathcal{D}'' \subseteq \mathcal{D}'$ be the set $\{R(h(\bar{c})) \mid R(\bar{c}) \in \mathcal{D}\}$. Clearly, $\langle \mathcal{D}'', \bar{b} \rangle$ is isomorphic to a database that can be obtained from $\langle \mathcal{D}, \bar{a} \rangle$ by identifying constants, i.e. isomorphic to a pointed database from $\text{id}(\mathcal{D}, \bar{a})$. Since \mathfrak{D} is closed under substructures, $\langle \mathcal{D}'', \bar{b} \rangle \in \mathfrak{D}$. By assumption, $h(\bar{a}) \in Q(\mathcal{D}'')$. Since $\mathcal{D}'' \rightarrow \mathcal{D}'$ and answers to queries from (FO, UCQ) are preserved under homomorphisms, $h(\bar{a}) \in Q(\mathcal{D}')$. Thus, $\bar{a} \in \text{app}_{\mathfrak{D}}^{\uparrow}(Q, \mathcal{D})$. \square

In the rest of this section, we only deal with Boolean (UCQ)s, so just write \mathcal{D} instead of $\langle \mathcal{D}, () \rangle$ and $\text{id}_{\mathfrak{D}}(\mathcal{D})$ instead of $\text{id}_{\mathfrak{D}}(\mathcal{D}, ())$. To prove Theorem 8, we show the following.

Lemma 25. *There is a boolean UCQ q such that the following problem is coNP-hard: Given a database \mathcal{D} , is it true that $\mathcal{D}' \models q$ for all $\mathcal{D}' \in \mathfrak{D}_1$ with $\mathcal{D} \rightarrow \mathcal{D}'$?*

Before giving a definition of the UCQ q , we describe the idea of the coNP-hardness proof. We reduce from the validity problem of propositional formulas in 3-DNF. Given a 3-DNF formula φ , we construct a database \mathcal{D}_{φ} such that $\mathcal{D}_{\varphi} \models q$ if and only if φ is valid. \mathcal{D}_{φ} consists of several variable gadgets \mathcal{D}_x , one for every variable x in φ . The variable gadget \mathcal{D}_x is a database that does not have treewidth 1, but there are precisely two intended ways to identify constants of \mathcal{D}_x to obtain a database of treewidth 1, where one way corresponds to assigning *true* to x , the other way to assigning *false*. There will be more than the two intended ways to identify constants of \mathcal{D}_x , but these unintended ways are detected using a UCQ $q_{\text{unintended}}$, which is

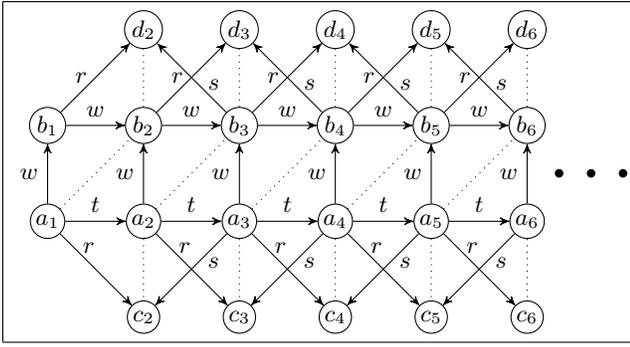


Figure 5: The figure shows a part of a variable gadget \mathcal{D}_x . The dotted lines correspond to the first of the two intended identifications, which stands for assigning *true* to x .

entailed in all databases that are obtained by an unintended way of identifying constants. However, if all variable gadgets are identified in an intended way, this corresponds to an assignment of the variables of φ . We then use a second UCQ q_{valid} to check whether there is a conjunct where all literals evaluate to *true*. The UCQ q is then defined as $q = q_{\text{unintended}} \vee q_{\text{valid}}$.

Let $\varphi = (\ell_1^1 \wedge \ell_1^2 \wedge \ell_1^3) \vee \dots \vee (\ell_n^1 \wedge \ell_n^2 \wedge \ell_n^3)$ be a propositional formula in 3-DNF, where every ℓ_i^j is either a variable or a negated variable. We start by constructing the variable gadgets. Let x be a variable of φ . The database \mathcal{D}_x takes the form of a grid of height 1 and width $N := 9 \cdot (2n - 1)$, plus some additional constants and facts that are used for technical reasons. We use constants of the form a_i^x and b_i^x , where $1 \leq i \leq N$, and of the form c_i^x and d_i^x where $1 < i < N$.

$$\begin{aligned} \mathcal{D}_x = & \{t(a_i^x, a_{i+1}^x) \mid 1 \leq i < N\} \cup \\ & \{w(b_i^x, b_{i+1}^x) \mid 1 \leq i < N\} \cup \\ & \{w(a_i^x, b_i^x) \mid 1 \leq i \leq N\} \cup \\ & \{r(a_i^x, c_{i+1}^x) \mid 1 \leq i \leq N - 2\} \cup \\ & \{s(a_{i+1}^x, c_i^x) \mid 2 < i \leq N - 1\} \cup \\ & \{r(b_i^x, d_{i+1}^x) \mid 1 \leq i \leq N - 2\} \cup \\ & \{s(b_{i+1}^x, d_i^x) \mid 2 < i \leq N - 1\} \end{aligned}$$

Figure 5 shows \mathcal{D}_x . So far, all \mathcal{D}_x are isomorphic. They will later be distinguished by adding certain unary facts. The two *intended identifications* in \mathcal{D}_x are the following:

1. Identify b_i^x with a_{i-1}^x for every i with $1 < i \leq N$. Furthermore, identify a_i^x with c_i^x and identify b_i^x with d_i^x for every i with $1 < i < N$. This identification corresponds to assigning *true* to x .
2. Identify b_i^x with a_{i+1}^x for every i with $1 \leq i < N$. Furthermore, identify a_i^x with c_i^x and identify b_i^x with d_i^x for every i with $1 < i < N$. This identification corresponds to assigning *false* to x .

We claim that the following UCQ detects precisely the unintended identifications:

$$q_{\text{unintended}} = (\exists x \exists y r(x, y) \wedge s(x, y)) \vee \exists x w(x, x) \vee \exists x t(x, x)$$

Lemma 26. Let \mathcal{D}_x be a variable gadget and $\mathcal{D}' \in \text{id}_{\mathfrak{D}_1}(\mathcal{D}_x)$. Then $\mathcal{D}' \models q_{\text{unintended}}$ if and only if the identification that was used to obtain \mathcal{D}' from \mathcal{D}_x is none of the two intended identifications.

Proof. One direction is easy to check: If \mathcal{D}' is obtained by one of the two intended identifications of constants, then $\mathcal{D}' \not\models q_{\text{unintended}}$. For the other direction, let $\mathcal{D}' \not\models q_{\text{unintended}}$. Since \mathcal{D}' has treewidth 1, it does not contain any cycle. For every i with $1 \leq i < N$, there is the cycle $w(a_i^x, b_i^x), w(b_i^x, b_{i+1}^x), w^-(b_{i+1}^x, a_{i+1}^x), t^-(a_{i+1}^x, a_i^x) \in \mathcal{D}_x$. In \mathcal{D}' , some of the four affected constants are identified. Since $\mathcal{D}' \not\models q_{\text{unintended}}$, no two constants that were adjacent in \mathcal{D}_x are identified, because this would have created a self-loop with the role w or t . The only two remaining possibilities are that (1) a_i^x was identified with b_{i+1}^x or that (2) a_{i+1}^x was identified with b_i^x . If for some i , a_i^x was identified with b_{i+1}^x , then b_{i+1}^x cannot be identified with a_{i+2}^x , since this would identify a_i^x with a_{i+2}^x , allowing a match for the first disjunct of $q_{\text{unintended}}$, mapping x to a_i^x and y to c_{i+1}^x . Thus, if a_i^x was identified with b_{i+1}^x for some i , then a_{i+1}^x was identified with b_{i+2}^x . Similarly, if b_i^x was identified with a_{i+1}^x for some i , then b_{i+1}^x was identified with a_{i+2}^x . By induction, this shows that either a_i^x is identified with b_{i+1}^x for all i with $1 \leq i < N$ or b_i^x is identified with a_{i+1}^x for all i with $1 \leq i < N$, just as in the two intended identifications.

It remains to show that for every i with $1 < i < N$, a_i^x is identified with c_i^x and b_i^x with d_i^x . Consider the cycle $t(a_{i-1}^x, a_i^x), t(a_i^x, a_{i+1}^x), s(a_{i+1}^x, c_i^x), r^-(c_i^x, a_{i-1}^x) \in \mathcal{D}_x$, that was collapsed in \mathcal{D}' . If a_{i-1}^x was identified with a_i^x or a_i^x with a_{i+1}^x , this would have created a match for $\exists x t(x, x)$. The only remaining options are identifying a_{i-1}^x with a_{i+1}^x or a_i^x with c_i^x . The first option would create a match for the first disjunct of $q_{\text{unintended}}$, so it follows that a_i^x was identified with c_i^x . A similar argument show that b_i^x is identified with d_i^x , considering the cycle $w(b_i^x, b_{i+1}^x), w(b_{i+1}^x, b_{i+2}^x), s(b_{i+2}^x, d_{i+1}^x), r^-(d_{i+1}^x, b_i^x) \in \mathcal{D}_x$. Thus, the identification of constants used to obtain \mathcal{D}' from \mathcal{D}_x is one of the two intended identifications. \square

Next, we construct \mathcal{D}_φ . We start by taking the disjoint union of all \mathcal{D}_x , for all variables x that appear in φ . Now we glue these gadgets together along their t -paths, i.e. for every i with $1 \leq i \leq N$ we do the following identification: identify all constants of the form a_i^x for some $x \in \text{var}(\varphi)$ to a single constant, which we just call a_i . Finally, we add some unary facts to encode φ in \mathcal{D} . For all $x \in \text{var}(\varphi)$ and all i, j such that $\ell_i^j \in \{x, \neg x\}$,

- add $T(b_k^x)$, where $k = 18(i - 1) + 3j - 1$;
- if $\ell_i^j = x$, add $A(a_k)$, where $k = 18(i - 1) + 3j - 2$;
- if $\ell_i^j = \neg x$, add $A(a_k)$, where $k = 18(i - 1) + 3j$.

It is clear that \mathcal{D}_φ can be constructed from φ in polynomial time. Note the effect that the two intended identifications have on \mathcal{D} : The identification that corresponds to an assignment that makes a literal ℓ_i^j become true creates a constant that satisfies both A and T .

Now we finish the definition of q by defining q_{valid} . Our aim is to build q_{valid} in such a way that it detects whether

the chosen assignment makes φ true, which is the case if and only if there is an i such that all ℓ_i^1 , ℓ_i^2 , and ℓ_i^3 evaluate to true under the assignment. This can be checked using a UCQ asking for certain t -paths that contain 3 different nodes that satisfy both A and T . For $S \in \{1, \dots, 9\}^3$, let $q_S = \exists x_1 \dots x_9 \bigwedge_{i=1}^8 t(x_i, x_{i+1}) \wedge \bigwedge_{i \in S} A(x_i) \wedge T(x_i)$, so q_S asks for a t -path of length 8, where all positions that appear in S are labelled with both A and T . Define

$$q_{\text{valid}} = \bigvee_{S \in \{1,3\} \times \{4,6\} \times \{7,9\}} q_S$$

and let $q = q_{\text{unintended}} \vee q_{\text{valid}}$. To finish the proof of Lemma 25, we need to argue the following:

Lemma 27. φ is valid if and only if for all $\mathcal{D}' \in \mathfrak{D}_1$ with $\mathcal{D}_\varphi \rightarrow \mathcal{D}'$, $\mathcal{D}' \models q$.

Proof. Let φ be valid. By Lemma 24, it is sufficient to show that $\mathcal{D}' \models q$ for all $\mathcal{D}' \in \text{id}_{\mathfrak{D}_1}(\mathcal{D}_\varphi)$. Let $\mathcal{D}' \in \text{id}_{\mathfrak{D}_1}(\mathcal{D}_\varphi)$ and assume, for the sake of contradiction, that $\mathcal{D}' \not\models q$. For every $x \in \text{var}(\varphi)$, \mathcal{D}_x is a substructure of \mathcal{D}_φ , so by Lemma 26, \mathcal{D}' was obtained from \mathcal{D}_φ by using one of the two intended identifications for every variable gadget \mathcal{D}_x . Recall that each of the intended identifications of \mathcal{D}_x corresponds to a truth assignment of x , so let $V : \text{var}(\varphi) \rightarrow \{0, 1\}$ be the corresponding assignment. Since φ is valid, there is one conjunct that evaluates to *true* under the assignment V , so there is an $i \in \{1, \dots, n\}$ such that all three literals ℓ_i^1 , ℓ_i^2 and ℓ_i^3 evaluate to true. Then q_{valid} has a match in \mathcal{D}' via the homomorphism $h(x_j) = a_{18(i-1)+j}$, a contradiction.

For the other direction, let φ be invalid, so let $V : \text{var}(\varphi) \rightarrow \{0, 1\}$ be an assignment that makes φ false. Let \mathcal{D}' be obtained from \mathcal{D}_φ by choosing for every substructure \mathcal{D}_x the identification corresponding to $V(x)$. It is easy to check that $\mathcal{D}' \in \mathfrak{D}_1$, in fact, \mathcal{D}' is a path with multi-edges. Since \mathcal{D}' was obtained by only using intended identifications and by Lemma 26, $\mathcal{D}' \not\models q_{\text{unintended}}$. For the sake of contradiction, assume that $\mathcal{D}' \models q_{\text{valid}}$, so there is a $(s_1, s_2, s_3) \in \{1, 3\} \times \{4, 6\} \times \{7, 9\}$ such that $\mathcal{D}' \models q_{(s_1, s_2, s_3)}$ via a homomorphism h . Since \mathcal{D}' contains only a single long t -path $t(a_1, a_2), \dots, t(a_{N-1}, a_N)$, the homomorphism h is of the form $h(x_i) = a_{k+1}$ for some $k \in \{1, \dots, N\}$. In particular, $A(a_{s_1}), A(a_{s_2}), A(a_{s_3}) \subseteq \mathcal{D}'$. However, \mathcal{D}_φ is constructed in such a way that whenever there are three occurrences $A(a_{k+s_1}), A(a_{k+s_2}), A(a_{k+s_3})$ within a radius of at most 8, then these were introduced for the three different literals ℓ_i^1, ℓ_i^2 and ℓ_i^3 of the i th conjunct for some $i \in \{1, \dots, n\}$. Furthermore, the existence of h also yields the three facts $T(a_{k+s_1}), T(a_{k+s_2}), T(a_{k+s_3}) \in \mathcal{D}'$. Since \mathcal{D}' was obtained from \mathcal{D}_φ using only intended identifications, and by construction of the \mathcal{D}_x , the fact $T(a_{s_j})$ can only be created by identifying a_{k+s_j} with $b_{k+s_j+1}^x$ or with $b_{k+s_j-1}^x$, where x is the variable of ℓ_i^j . But these identifications correspond to an assignment, that makes all three literals ℓ_i^1, ℓ_i^2 and ℓ_i^3 evaluate to true, contradicting the assumption that V makes φ false. \square

We describe how the construction can be adapted slightly to show that \mathfrak{D}_1 -database strengthening approximation in $(\mathcal{EL}, \text{CQ})$ is CONP-hard as well. Let $q_0 = \exists x \exists y r(x, y) \wedge s(x, y) \wedge M(y)$. The idea is to design an ontology \mathcal{O} and database \mathcal{D} in such a way that for every $\mathcal{D}' \in \text{id}_{\mathfrak{D}_1}(\mathcal{D})$ we have $\mathcal{D}' \models q$ if and only if $\mathcal{O}, \mathcal{D}' \models q_0$.

In the construction of \mathcal{D}_x , we add the facts $M(c_i^x)$ and $M(d_i^x)$ for all i with $1 < i < N$. Furthermore, we introduce constants e_i for $1 \leq i \leq N$ and add the facts $r(a_i, e_i), s(a_i, e_i)$ and $v(e_i, a_i)$. We also introduce new unary symbols A_1, A_2, A_3 and B_1, B_2, B_3 and add $A_{i \bmod 3}(a_i^x)$ and $B_{i \bmod 3}(b_i)$ for all $x \in \text{var}(\varphi)$ and i with $1 \leq i \leq N$.

The ontology is used to translate every match of q into a match of q_0 . Recall that

$$q_{\text{unintended}} = (\exists x \exists y r(x, y) \wedge s(x, y)) \vee \exists x v(x, x) \vee \exists x t(x, x).$$

For the first disjunct of $q_{\text{unintended}}$, there is nothing to do, since we already added the facts $M(c_i^x)$ and $M(d_i^x)$ to \mathcal{D} . For the second and third disjunct, we introduce the concept inclusions $\exists v.C \sqsubseteq M$ for every $C \in \{A_1 \sqcap A_2, A_1 \sqcap A_3, A_2 \sqcap A_3, B_1 \sqcap B_2, B_1 \sqcap B_3, B_2 \sqcap B_3, A_1 \sqcap B_1, A_2 \sqcap B_2, A_3 \sqcap B_3\}$. Every disjunct q_S of q_{valid} can be seen as an \mathcal{EL} -concept C_S . We introduce a concept inclusion $\exists v.C_S \sqsubseteq M$ for each such C_S . It can then be verified that $\mathcal{O}, \mathcal{D} \models q_0$ if and only if φ is valid.

It is an interesting question whether the lower bound can already be established for $(\mathcal{EL}, \text{AQ})$, or for the empty ontology with CQs or AQs. We leave this as an open problem.