Error-Tolerant Reasoning in the Description Logic $\mathcal{EL}$ Based on Optimal Repairs*

Franz Baader, Francesco Kriegel, and Adrian Nuradiansyah

Theoretical Computer Science, Technische Universität Dresden, Dresden, Germany
firstname.lastname@tu-dresden.de

Abstract. Ontologies based on Description Logic (DL) represent general background knowledge in a terminology (TBox) and the actual data in an ABox. Both human-made and machine-learned data sets may contain errors, which are usually detected when the DL reasoner returns unintuitive or obviously incorrect answers to queries. To eliminate such errors, classical repair approaches offer as repairs maximal subsets of the ABox not having the unwanted answers w.r.t. the TBox. It is, however, not always clear which of these classical repairs to use as the new, corrected data set. Error-tolerant semantics instead takes all repairs into account: cautious reasoning returns the answers that follow from all classical repairs whereas brave reasoning returns the answers that follow from some classical repair. It is inspired by inconsistency-tolerant reasoning and has been investigated for the DL $\mathcal{EL}$, but in a setting where the TBox rather than the ABox is repaired. In a series of papers, we have developed a repair approach for ABoxes that improves on classical repairs in that it preserves a maximal set of consequences (i.e., answers to queries) rather than a maximal set of ABox assertions. The repairs obtained by this approach are called optimal repairs. In the present paper, we investigate error-tolerant reasoning in the DL $\mathcal{EL}$, but we repair the ABox and use optimal repairs rather than classical repairs as the underlying set of repairs. To be more precise, we consider a static $\mathcal{EL}$ TBox (which is assumed to be correct), represent the data by a quantified ABox (where some individuals may be anonymous), and use $\mathcal{EL}$ concepts as queries (instance queries). We show that brave entailment of instance queries can be decided in polynomial time. Cautious entailment can be decided by a coNP procedure, but is still in P if the TBox is empty.

1 Introduction

Description Logics (DLs) [2] are a prominent family of logic-based knowledge representation formalisms, which offer a good compromise between expressiveness and the complexity of reasoning and are the formal basis for the Web ontology language OWL.[1] Here we concentrate on the inexpressive and tractable DL

* partially supported by the AI competence center ScaDS.AI Dresden/Leipzig and the German Research Foundation (DFG) in Project 430150274 and SFB/TRR 248.

[1] https://www.w3.org/TR/owl2-overview/

$\mathcal{EL}$ [1], which is frequently used to represent ontologies in biology and medicine, such as the large medical ontology SNOMED CT.[2]

Like all large human-made digital artefacts, the ontologies employed in such applications may contain errors, and this problem gets even worse if parts of the ontology (usually the data) are automatically generated by inexact methods based on information retrieval or machine learning. Errors are often detected when reasoning finds an inconsistency or generates unintuitive consequences. To correct such a mistake, classical repair approaches propose to use maximal subsets of the ontology as repairs [9, 17, 19]. While these approaches preserve as many of the original axioms as possible, they may not be optimal w.r.t. preserving consequences. In a series of papers [3, 7, 8], we have investigated how to characterize and compute optimal repairs, which are defined to be ontologies entailed by the erroneous ontology whose consequence sets are maximal among all such ontologies. To illustrate the difference between classical and optimal repairs, assume that the (quantified) ABox consists of the assertions $\mathit{owns}(\mathit{Ralf}, x)$, $\mathit{Red}(x)$, and $\mathit{Bike}(x)$, where x is an anonymous individual, but that the consequence $\exists\, \mathit{owns}.(\mathit{Red} \sqcap \mathit{Bike})(\mathit{Ralf})$ is assumed to be incorrect. There are three classical repairs, obtained by respectively removing one of the assertions, but only one optimal repair, which consists of the assertions $\mathit{owns}(\mathit{Ralf}, y)$, $\mathit{Red}(y)$, $\mathit{owns}(\mathit{Ralf}, z)$, and $\mathit{Bike}(z)$. Clearly, this repair preserves more consequences (in the sense of instance relationships for Ralf) than each of the classical repairs.

In general, a given repair problem may have exponentially many repairs, both in the classical and the optimal sense, and it is often hard to decide which one to use. Error-tolerant reasoning does not commit to a single repair, but rather reasons w.r.t. all of them (within the classical or the optimal setting): cautious reasoning returns the answers that follow from all repairs whereas brave reasoning returns the answers that follow from some repair. For classical repairs of TBoxes in $\mathcal{EL}$, it was first investigated in [16, 18], where it was shown that brave entailment is NP-complete and cautious entailment is coNP-complete. For more expressive DLs that can create inconsistencies, error-tolerant reasoning had been considered before, for the case where the error is an inconsistency, under the name of inconsistency-tolerant reasoning [10, 11, 15]. This latter work also uses the classical notion of repair.

In the present paper, we investigate error-tolerant reasoning in the DL $\mathcal{EL}$, using optimal repairs of ABoxes as the underlying set of repairs. To be more precise, we consider a static $\mathcal{EL}$ TBox (which is assumed to be correct, and thus cannot be changed), represent the data by a quantified ABox, and consider instance relationships between individuals and $\mathcal{EL}$ concepts as relevant consequences. In [3] it is shown that, in this setting, each repair is entailed by an optimal repair and that every optimal repair is equivalent to a so-called canonical repair, which is induced by a polynomially large repair seed function.[3]

[2] `https://www.snomed.org/`

[3] Since we are only interested in instance relationships, the appropriate entailment and equivalence relations between quantified ABoxes are IQ-entailment and IQ-equivalence [3].

For the case of brave reasoning, it is actually sufficient to know that every repair is entailed by an optimal one. From this, we obtain that a set of concept assertions is entailed by some optimal repair if, and only if, it is itself a repair. The latter property can be tested by performing a polynomial number of polynomial-time instance tests, which shows that brave reasoning is tractable.

Dealing with cautious reasoning is more complicated since we really need to check what is entailed by all optimal repairs. The first problem is then that, while seed functions are of polynomial size, the canonical repairs they induce may be of exponential size. The solution to this problem is that we work directly with the seed functions without computing the induced repairs. This is possible since we can show that entailment from the canonical repair induced by a given seed function can actually be decided in polynomial time in the size of the seed function, and not just in the size of the repair. The second problem is that, while the set of canonical repairs contains (up to equivalence) all optimal repairs, there may exist non-optimal canonical repairs. Thus, cautious reasoning cannot be done w.r.t. all canonical repairs. Nevertheless, for cautious entailment w.r.t. the empty TBox, we are able to provide a direct characterization, and can show that this condition can be checked in polynomial time. For cautious reasoning w.r.t. a non-empty TBox, we use the fact (shown in [6]) that the optimal repairs are induced by seed functions that are minimal w.r.t. an appropriate pre-order on such functions. Non-entailment can then be tested using a guess-and-check approach that guesses a seed function, checks whether it is minimal, and then checks non-entailment. To show that this yields an NP procedure for non-entailment, we must prove that minimality of a seed function can be tested in polynomial time. Overall, we obtain tractability of cautious reasoning w.r.t. the empty TBox, and a coNP upper bound for cautious reasoning w.r.t. a non-empty TBox. Whether this bound is tight remains an open problem.

2 Preliminaries

We start with recalling the DL $\mathcal{EL}$ as well as $\mathcal{EL}$ TBoxes and ABoxes, and then introduce quantified ABoxes and the entailment relation used in this paper to compare them. We assume that the reader is familiar with the basic notions of description and first-order logic and base our presentation on the one in [5].

The Description Logic $\mathcal{EL}$. Starting from a signature Σ, which is a disjoint union of a set Σ_I of *individual names*, a set Σ_C of *concept names*, and a set Σ_R of *role names*, $\mathcal{EL}$ *concept descriptions* are built using the grammar $C ::= \top \mid A \mid C \sqcap C \mid \exists r.\, C$, where A ranges over Σ_C and r over Σ_R. An $\mathcal{EL}$ *concept assertion* is of the form $C(a)$ where C is an $\mathcal{EL}$ concept description and $a \in \Sigma_\mathsf{I}$, a *role assertion* is of the form $r(a, b)$ where $r \in \Sigma_\mathsf{R}$ and $a, b \in \Sigma_\mathsf{I}$, and an $\mathcal{EL}$ *concept inclusion (CI)* is of the form $C \sqsubseteq D$ for concept descriptions C, D. An $\mathcal{EL}$ *ABox* $\mathcal{A}$ is a finite set of concept assertions and role assertions, and an $\mathcal{EL}$ *TBox* $\mathcal{T}$ is a finite set of concept inclusions. Since $\mathcal{EL}$ is the only DL considered in this paper, we will sometimes omit the prefix "$\mathcal{EL}$," and we will use "concept" as an abbreviation for "concept description."

The semantics of $\mathcal{EL}$ can be defined either directly in a model-theoretic way of by a translation into first-order logic (FO) [2]. In the translation, the elements of Σ_I, Σ_C, and Σ_R are respectively viewed as constant symbols, unary predicate symbols, and binary predicate symbols. $\mathcal{EL}$ concepts C are inductively translated into FO formulas $\phi_C(x)$ with one free variable x:

- concept A for $A \in \Sigma_\mathsf{C}$ is translated into $A(x)$ and $\top$ into $A(x) \vee \neg A(x)$ for an arbitrary $A \in \Sigma_\mathsf{C}$;
- if C, D are translated into $\phi_C(x)$ and $\phi_D(x)$, then $C \sqcap D$ is translated into $\phi_C(x) \wedge \phi_D(x)$ and $\exists r. C$ into $\exists y. (r(x, y) \wedge \phi_C(y))$, where $\phi_C(y)$ is obtained from $\phi_C(x)$ by replacing the free variable x by a different variable y.

CIs $C \sqsubseteq D$ are translated into sentences $\phi_{C \sqsubseteq D} := \forall x. (\phi_C(x) \to \phi_D(x))$ and TBoxes $\mathcal{T}$ into $\phi_\mathcal{T} := \bigwedge_{C \sqsubseteq D \in \mathcal{T}} \phi_{C \sqsubseteq D}$. Concept assertions $C(a)$ are translated into $\phi_C(a)$, role assertions $r(a, b)$ stay the same, and ABoxes $\mathcal{A}$ are translated into the conjunction $\phi_\mathcal{A}$ of the translations of their assertions.

Let α, β be ABoxes, concept inclusions, or concept assertions (possibly not both of the same kind), and $\mathcal{T}$ an $\mathcal{EL}$ TBox. Then we say that α *entails* β *w.r.t.* $\mathcal{T}$ (written $\alpha \models^\mathcal{T} \beta$) if the implication $(\phi_\alpha \wedge \phi_\mathcal{T}) \to \phi_\beta$ is valid according to the semantics of FO. Furthermore, α and β are *equivalent w.r.t.* $\mathcal{T}$ (written $\alpha \equiv^\mathcal{T} \beta$), if $\alpha \models^\mathcal{T} \beta$ and $\beta \models^\mathcal{T} \alpha$. In case $\mathcal{T} = \emptyset$, we will sometimes write $\models$ instead of $\models^\emptyset$. If $\emptyset \models^\mathcal{T} C \sqsubseteq D$, then we also write $C \sqsubseteq^\mathcal{T} D$ and say that C *is subsumed by* D *w.r.t.* $\mathcal{T}$; in case $\mathcal{T} = \emptyset$ we simply say that C is subsumed by D. The subsumption problem in $\mathcal{EL}$ is known to be decidable in polynomial time [1], and the same is true for all the entailment problems introduced above.

Quantified ABoxes. A *quantified ABox* (qABox) $\exists X. \mathcal{A}$ consists of a set X of *variables*, which is disjoint with Σ, and a *matrix* $\mathcal{A}$, which is a finite set of concept assertions $A(u)$ and role assertions $r(u, v)$, where $A \in \Sigma_\mathsf{C}$, $r \in \Sigma_\mathsf{R}$ and $u, v \in \Sigma_\mathsf{I} \cup X$. The matrix is an ABox built over the extended signature $\Sigma \cup X$, but cannot contain complex concept descriptions. An *object* of $\exists X. \mathcal{A}$ is either an individual name in Σ_I or a variable in X.

Like $\mathcal{EL}$ ABoxes, quantified ABox $\exists X. \mathcal{A}$ can be translated into FO sentences, but where the elements of X are viewed as first-order variables rather than constants and are existentially quantified. Thus, entailment between two qABoxes (written $\exists X. \mathcal{A} \models^\mathcal{T} \exists Y. \mathcal{B}$) and between a qABox and a concept assertion (written $\exists X. \mathcal{A} \models^\mathcal{T} C(a)$) w.r.t. a TBox $\mathcal{T}$ can again be defined using the semantics of first-order logic.[4] If $\exists X. \mathcal{A} \models^\mathcal{T} C(a)$, then a is called an *instance* of C w.r.t. $\exists X. \mathcal{A}$ and $\mathcal{T}$.

Syntactically, not every $\mathcal{EL}$ ABox is a qABox, since $\mathcal{EL}$ ABoxes may contain concept assertions $C(a)$ for complex concepts C. However, every $\mathcal{EL}$ ABox can be translated into an equivalent qABox, by writing the FO translation of

[4] see [3, 8] for more information on qABoxes

complex concepts C as a qABox. For example, if $C = \exists r.(A \sqcap B)$, then the $\mathcal{EL}$ ABox $\{C(a)\}$ is equivalent to the qABox $\exists\{x\}.\{r(a,x), A(x), B(x)\}$. Conversely, not every qABox can be expressed by an $\mathcal{EL}$ ABox since qABoxes may contain cyclic role relations between variables. For example, if $\mathcal{T}$ is empty, then the qABox $\exists\{x\}.\{r(a,x), r(x,x)\}$ is not equivalent to an $\mathcal{EL}$ ABox [5]. One might be tempted to think that one can just forget about the existential quantifier and use an individual b instead of the variable x. However, the ABox $\{r(a,b), r(b,b)\}$ is not equivalent to the above qABox since it entails non-trivial instance relationships for b, whereas $\exists\{x\}.\{r(a,x), r(x,x)\}$ does not. Also note that, while entailment between $\mathcal{EL}$ ABoxes and entailment of a concept assertion by a qABox can be decided in polynomial time, the entailment problem between qABoxes is NP-complete [3, 8].

However, since in this paper we are only interested in the instance relationships that a given qABox entails, we can restrict our attention to IQ-entailment between qABoxes: the qABox $\exists X.\mathcal{A}$ IQ-*entails* the qABox $\exists Y.\mathcal{B}$ w.r.t. $\mathcal{T}$ (written $\exists X.\mathcal{A} \models^{\mathcal{T}}_{\mathsf{IQ}} \exists Y.\mathcal{B}$) if $\exists Y.\mathcal{B} \models^{\mathcal{T}} C(a)$ implies $\exists X.\mathcal{A} \models^{\mathcal{T}} C(a)$ for each $\mathcal{EL}$ concept assertion $C(a)$. In contrast to the FO entailment introduced above, IQ-entailment between qABoxes can be decided in polynomial time. This is a consequence of the following result from [3]: given a qABox $\exists X.\mathcal{A}$ and an $\mathcal{EL}$ TBox $\mathcal{T}$, one can compute in polynomial time an IQ-*saturation* $\mathsf{sat}^{\mathcal{T}}_{\mathsf{IQ}}(\exists X.\mathcal{A})$ such that the following statements are equivalent:

- $\exists X.\mathcal{A} \models^{\mathcal{T}}_{\mathsf{IQ}} \exists Y.\mathcal{B}$
- $\mathsf{sat}^{\mathcal{T}}_{\mathsf{IQ}}(\exists X.\mathcal{A}) \models_{\mathsf{IQ}} \exists Y.\mathcal{B}$
- There is a simulation from $\exists Y.\mathcal{B}$ to $\mathsf{sat}^{\mathcal{T}}_{\mathsf{IQ}}(\exists X.\mathcal{A})$.

The notion of simulation employed here is the usual one for labeled graphs, whose existence can be decided in polynomial time (see [3] for details).

3 Optimal and Canonical Repairs

We first introduce the notion of an optimal repair w.r.t. IQ-entailment and recall the approach for obtaining canonical IQ-repairs based on repair seed functions described in [3]. Then, we show that reasoning w.r.t. canonical repairs can be performed by considering the seed function rather than the induced canonical repair. Since the optimal repairs are exactly the canonical ones induced by minimal seed functions [3], we also investigate how minimality of a seed function can be decided. The reason for employing IQ-entailment is that we are only interested in the instance relationships entailed by a given qABox and TBox. We use repair requests to indicate which consequences are considered to be erroneous, and thus need to be removed. Formally, a *repair request* $\mathcal{R}$ is a finite set of concept assertions.

Definition 1. *Let $\mathcal{T}$ be an $\mathcal{EL}$ TBox, $\exists X.\mathcal{A}$ a qABox, and $\mathcal{R}$ a repair request.*

- *The qABox $\exists Y.\mathcal{B}$ is an* IQ-repair *of $\exists X.\mathcal{A}$ for $\mathcal{R}$ w.r.t. $\mathcal{T}$ if $\exists X.\mathcal{A} \models^{\mathcal{T}}_{\mathsf{IQ}} \exists Y.\mathcal{B}$ and $\exists Y.\mathcal{B} \not\models^{\mathcal{T}} C(a)$ for each $C(a) \in \mathcal{R}$.*

– *Such a repair* $\exists Y.\mathcal{B}$ *is* optimal *if there is no* IQ*-repair* $\exists Z.\mathcal{C}$ *such that* $\exists Z.\mathcal{C} \models^{\mathcal{T}}_{\mathsf{IQ}} \exists Y.\mathcal{B}$*, but* $\exists Y.\mathcal{B} \not\models^{\mathcal{T}}_{\mathsf{IQ}} \exists Z.\mathcal{C}$*.*

Not every repair request has a repair, but the ones that have can easily be identified. We call a repair request $\mathcal{R}$ *solvable* w.r.t. a TBox $\mathcal{T}$ if, for each quantified ABox $\exists X.\mathcal{A}$, there exists a repair of $\exists X.\mathcal{A}$ for $\mathcal{R}$ w.r.t. $\mathcal{T}$. As mentioned in [3], this is the case iff $\top \not\sqsubseteq^{\mathcal{T}} C$ for each $C(a) \in \mathcal{R}$.

In general, a given repair instance $\mathcal{T}, \exists X.\mathcal{A}, \mathcal{R}$ may have exponentially many non-equivalent optimal repairs. Repair seed functions can be used to define (a superset of) these repairs, by specifying, for each individual a in $\mathcal{A}$, which atoms should not hold for a in the repair. To take the TBox into account, one first constructs the IQ-saturation $\exists Y.\mathcal{B} := \mathsf{sat}^{\mathcal{T}}_{\mathsf{IQ}}(\exists X.\mathcal{A})$. We denote the set of all subconcepts of concepts occurring in $\mathcal{R}$ or $\mathcal{T}$ with $\mathsf{Sub}(\mathcal{R}, \mathcal{T})$. An *atom* is either a concept name or an existential restriction, and we denote the set of atoms in $\mathsf{Sub}(\mathcal{R}, \mathcal{T})$ with $\mathsf{Atoms}(\mathcal{R}, \mathcal{T})$.

Definition 2. *Let* $\mathcal{T}$ *be an* $\mathcal{EL}$ *TBox,* $\exists X.\mathcal{A}$ *a qABox,* $\mathcal{R}$ *a repair request, and* $\exists Y.\mathcal{B}$ *the* IQ*-saturation of* $\exists X.\mathcal{A}$ *w.r.t.* $\mathcal{T}$*. A* repair type *for an object* u *of* $\exists Y.\mathcal{B}$ *is a subset* $\mathcal{K}$ *of* $\mathsf{Atoms}(\mathcal{R}, \mathcal{T})$ *that satisfies the following three conditions:*

1. $K \not\sqsubseteq^{\emptyset} K'$ *for all distinct atoms* $K, K' \in \mathcal{K}$*.*
2. $\mathcal{B} \models K(u)$ *for every atom* $K \in \mathcal{K}$*.*
3. $\mathcal{K}$ *is* premise-saturated*, i.e., if* $K \in \mathcal{K}$ *and* $C \in \mathsf{Sub}(\mathcal{R}, \mathcal{T})$ *are such that* $\mathcal{B} \models C(u)$ *and* $C \sqsubseteq^{\mathcal{T}} K$*, then there is* $K' \in \mathcal{K}$ *with* $C \sqsubseteq^{\emptyset} K'$*.*[5]

A repair seed function (rsf) s *assigns to each individual name* $a \in \Sigma_{\mathsf{I}}$ *a repair type* $s(a)$ *such that, for each unwanted consequence* $C(a) \in \mathcal{R}$ *with* $\mathcal{B} \models C(a)$*, there is an atom* $K \in s(a)$ *with* $C \sqsubseteq^{\emptyset} K$*.*

As shown in [3], each rsf s induces a *canonical* IQ-repair, denoted as $\mathsf{rep}^{\mathcal{T}}_{\mathsf{IQ}}(\exists X.\mathcal{A}, s)$, and the set of canonical IQ-repairs covers all IQ-repairs in the sense that every repair is IQ-entailed by a canonical one. In particular, this implies that, up to IQ-equivalence, the set of canonical IQ-repairs contains all optimal IQ-repairs, and the set of optimal IQ-repairs also covers all IQ-repairs.

For the purposes of this paper, the exact definition of $\mathsf{rep}^{\mathcal{T}}_{\mathsf{IQ}}(\exists X.\mathcal{A}, s)$ is not relevant since we intend to work directly with the (polynomial-sized) rsf s rather than the (exponentially large) induced canonical repair. An important result that helps us to do this is the following lemma, which is an extension of Lemma XII in [4], whose proof is similar to the proof of Lemma VI in [14].

Lemma 3. *Let* s *be a repair seed function,* b *an individual in* $\mathcal{A}$*, and* C *an* $\mathcal{EL}$ *concept. Then* $\mathsf{rep}^{\mathcal{T}}_{\mathsf{IQ}}(\exists X.\mathcal{A}, s) \models^{\mathcal{T}} C(b)$ *iff* $\exists X.\mathcal{A} \models^{\mathcal{T}} C(b)$ *and* $s(b)$ *does not contain an atom that subsumes* C *w.r.t.* $\mathcal{T}$*.*

Since the right-hand side of this equivalence can obviously be checked in polynomial time (since $\exists X.\mathcal{A} \models^{\mathcal{T}} C(b)$ iff $\mathcal{A} \models^{\mathcal{T}} C(b)$) and $s(b)$ is of polynomial size, we obtain the following complexity result.

[5] A *repair pre-type* need only satisfy the first two conditions. If the TBox is empty, then this third condition is trivially true since one can take $K' = K$.

Proposition 4. *Given a qABox $\exists X.\mathcal{A}$, an $\mathcal{EL}$ TBox $\mathcal{T}$, a repair request $\mathcal{R}$, a repair seed function s, and an $\mathcal{EL}$ concept assertion $C(b)$, we can decide in polynomial time (in the size of $\exists X.\mathcal{A}$, $\mathcal{T}$, and $\mathcal{R}$) whether $C(b)$ is entailed w.r.t. $\mathcal{T}$ by the canonical* IQ*-repair induced by s.*

The set of canonical repairs may contain non-optimal repairs. A simple example is given by the empty TBox, the qABox $\exists\emptyset.\{A(a), B(a)\}$, and the repair request $\mathcal{R} = \{(A \sqcap B)(a)\}$. There are three seed functions s_1, s_2, s_3 with $s_1(a) = \{A\}, s_2(a) = \{B\}, s_3(a) = \{A, B\}$, which respectively induce the canonical repairs $\exists\emptyset.\{B(a)\}$, $\exists\emptyset.\{A(a)\}$, and $\exists\emptyset.\emptyset$. Whereas the first two are optimal repairs, the latter one is not optimal; in fact, it is strictly entailed by each of the former ones. Obviously, the reason for this is that $s_3(a)$ is contained both in $s_1(a)$ and in $s_2(a)$.

More generally, we can reflect entailment between canonical repairs by the following covering relation between seed functions. Given sets $\mathcal{K}$ and $\mathcal{L}$ of concept descriptions, we say that $\mathcal{K}$ is *covered* by $\mathcal{L}$ (written $\mathcal{K} \leq \mathcal{L}$) if, for each $K \in \mathcal{K}$, there is $L \in \mathcal{L}$ such that $K \sqsubseteq^\emptyset L$. Applying the covering relation argumentwise yields the following pre-order on seed functions: $s \leq t$ if $s(a) \leq t(a)$ for each $a \in \Sigma_\mathsf{I}$. The following result, which is an easy consequence of Lemma 3, was already mentioned in [6].

Lemma 5. $s \leq t$ *iff* $\mathsf{rep}^{\mathcal{T}}_{\mathsf{IQ}}(\exists X.\mathcal{A}, s) \models^{\mathcal{T}}_{\mathsf{IQ}} \mathsf{rep}^{\mathcal{T}}_{\mathsf{IQ}}(\exists X.\mathcal{A}, t)$.

Given any pre-order $\leq$, we write $\alpha < \beta$ if $\alpha \leq \beta$ and $\beta \not\leq \alpha$, and say that α is $\leq$-*minimal* ($\leq$-*maximal*) if there is no β such that $\beta < \alpha$ ($\alpha < \beta$). For repair seed functions s, t we have $s < t$ iff $s(a) \leq t(a)$ for all $a \in \Sigma_\mathsf{I}$ and there is $b \in \Sigma_\mathsf{I}$ with $s(b) < t(b)$. As an immediate consequence of Lemma 5, we obtain that the optimal repairs are induced by the minimal seed functions.

Proposition 6. *If s is a $\leq$-minimal rsf, then $\mathsf{rep}^{\mathcal{T}}_{\mathsf{IQ}}(\exists X.\mathcal{A}, s)$ is an optimal* IQ*-repair, and every optimal* IQ*-repair is* IQ*-equivalent to a canonical repair $\mathsf{rep}^{\mathcal{T}}_{\mathsf{IQ}}(\exists X.\mathcal{A}, s)$ for a $\leq$-minimal rsf s.*

In the rest of this section we show that $\leq$-minimality of seed functions can be decided in polynomial time. More precisely, we characterise non-minimality by showing how, for a given repair type, we can decide whether there exists a repair type that is strictly covered by it. As before, we denote by $\exists Y.\mathcal{B}$ the saturation $\mathsf{sat}^{\mathcal{T}}_{\mathsf{IQ}}(\exists X.\mathcal{A})$. We start by showing how, for a given repair type $\mathcal{K}$, a non-empty set $\mathcal{M}$ of atoms covered by it can be employed to construct a repair pre-type that is strictly covered by $\mathcal{K}$.

Definition 7. *Let $\mathcal{K}$ be a repair type for u and $\mathcal{M}$ be a non-empty subset of $\mathsf{Atoms}(\mathcal{R}, \mathcal{T})$ such that $\mathcal{M} \leq \mathcal{K}$. We define the* lowering *of $\mathcal{K}$ w.r.t. $\mathcal{M}$ by*

$$\mathsf{low}(\mathcal{K}, \mathcal{M}) := \mathsf{Max}_{\sqsubseteq^\emptyset} \left\{ E \;\middle|\; \begin{array}{l} E \in \mathsf{Atoms}(\mathcal{R}, \mathcal{T}),\ \mathcal{B} \models E(u), \\ E \sqsubseteq^\emptyset K \text{ for some } K \in \mathcal{K}, \\ M \not\sqsubseteq^\emptyset E \text{ for each } M \in \mathcal{M} \end{array} \right\}.$$

Due to the $\mathsf{Max}_{\sqsubseteq^\emptyset}$ operator, which selects a representative for each equivalence class of $\sqsubseteq^\emptyset$-maximal elements, and the condition that each atom E in $\mathsf{low}(\mathcal{K}, \mathcal{M})$ must satisfy $\mathcal{B} \models E(u)$, we know that $\mathsf{low}(\mathcal{K}, \mathcal{M})$ is a repair pre-type for u. Next, we show that $\mathcal{K}$ strictly covers $\mathsf{low}(\mathcal{K}, \mathcal{M})$.

Lemma 8. $\mathsf{low}(\mathcal{K}, \mathcal{M}) < \mathcal{K}$

Proof. By definition, each atom in $\mathsf{low}(\mathcal{K}, \mathcal{M})$ is subsumed by some atom in $\mathcal{K}$, which means that $\mathsf{low}(\mathcal{K}, \mathcal{M}) \leq \mathcal{K}$.

To show that $\mathcal{K} \not\leq \mathsf{low}(\mathcal{K}, \mathcal{M})$, we consider an element $M \in \mathcal{M}$, which exists since we have assumed $\mathcal{M} \neq \emptyset$. Since $\mathcal{M} \leq \mathcal{K}$, there is an atom K in $\mathcal{K}$ such that $M \sqsubseteq^\emptyset K$. We show that K is not subsumed by any atom in $\mathsf{low}(\mathcal{K}, \mathcal{M})$.

Assume to the contrary that $K \sqsubseteq^\emptyset E$ for some atom $E \in \mathsf{low}(\mathcal{K}, \mathcal{M})$. Then $E \sqsubseteq^\emptyset K'$ for some $K' \in \mathcal{K}$, and thus $K \sqsubseteq^\emptyset K'$. Since the repair type $\mathcal{K}$ cannot contain distinct $\sqsubseteq^\emptyset$-comparable atoms, K and K' must be equal. We infer from $K \sqsubseteq^\emptyset E \sqsubseteq^\emptyset K'$ that E and K are equivalent, and thus $M \sqsubseteq^\emptyset K$ yield $M \sqsubseteq^\emptyset E$. This contradicts our assumption that $E \in \mathsf{low}(\mathcal{K}, \mathcal{M})$ □

The lowering of $\mathcal{K}$ w.r.t. $\mathcal{M}$ need not be a repair type, but we can construct, for each atom $D \in \mathcal{K}$, a set $\mathcal{M}_D$ such that $\mathsf{low}(\mathcal{K}, \mathcal{M}_D)$ is a repair type.

Definition 9. *Let $\mathcal{K}$ be a repair type for u and $D \in \mathcal{K}$. We inductively define the following sets:*

$$\mathcal{M}_D^0 \coloneqq \{D\}$$

$$\mathcal{M}_D^{i+1} \coloneqq \mathcal{M}_D^i \cup \left\{ F \;\middle|\; \begin{array}{l} F \in \mathsf{low}(\mathcal{K}, \mathcal{M}_D^i) \textit{ and there is } C \in \mathsf{Sub}(\mathcal{R}, \mathcal{T}) \\ \textit{such that } \mathcal{B} \models C(u),\ C \sqsubseteq^\mathcal{T} F,\ \{C\} \not\leq \mathsf{low}(\mathcal{K}, \mathcal{M}_D^i) \end{array} \right\}$$

We further set $\mathcal{M}_D \coloneqq \mathcal{M}_D^j$ where j is the minimal index such that $\mathcal{M}_D^{j+1} = \mathcal{M}_D^j$.

Since we can show by induction that $\mathcal{M}_D^i$ is non-empty and covered by $\mathcal{K}$ for all $i \geq 0$, $\mathcal{K}$ and $\mathcal{M}_D^i$ satisfy the conditions of Definition 7 on the arguments of low in the definition of $\mathcal{M}_D^{i+1}$.

Lemma 10. $\mathsf{low}(\mathcal{K}, \mathcal{M}_D)$ *is a repair type for u.*

Proof. We have already seen that $\mathsf{low}(\mathcal{K}, \mathcal{M}_D)$ is a repair pre-type. It remains to prove that it is premise-saturated. Thus, let $F \in \mathsf{low}(\mathcal{K}, \mathcal{M}_D)$ and $C \in \mathsf{Sub}(\mathcal{R}, \mathcal{T})$ be such that $\mathcal{B} \models C(u)$ and $C \sqsubseteq^\mathcal{T} F$, and assume that C is not subsumed by any atom in $\mathsf{low}(\mathcal{K}, \mathcal{M}_D) = \mathsf{low}(\mathcal{K}, \mathcal{M}_D^j)$. Then $F \in \mathcal{M}_D^{j+1} = \mathcal{M}_D$, which yields a contradiction since $F \in \mathsf{low}(\mathcal{K}, \mathcal{M}_D)$ requires that F does not subsume any atom in $\mathcal{M}_D$. □

Next, we characterize the repair types that are strictly covered by a given repair type $\mathcal{K}$.

Lemma 11. *Let $\mathcal{K}$ and $\mathcal{L}$ be repair types for u. Then, $\mathcal{L} < \mathcal{K}$ iff there is some $D \in \mathcal{K}$ such that $\mathcal{L} \leq \mathsf{low}(\mathcal{K}, \mathcal{M}_D)$.*

Proof. The if direction follows directly from Lemma 8. To show the only-if direction, assume that $\mathcal{L} < \mathcal{K}$, i.e., $\mathcal{L} \leq \mathcal{K}$ and $\mathcal{K} \not\leq \mathcal{L}$. The latter yields an atom $D \in \mathcal{K}$ that is not subsumed by any atom in $\mathcal{L}$. We show by induction that $\mathcal{L} \leq \mathsf{low}(\mathcal{K}, \mathcal{M}_D^i)$ for all $i \geq 0$.

In the base case ($i = 0$), we have $\mathcal{M}_D^0 = \{D\}$. Consider an atom $L \in \mathcal{L}$. Since $\mathcal{L}$ is a repair type for u, it holds that $\mathcal{B} \models L(u)$. Since $\mathcal{L} \leq \mathcal{K}$, there is an atom $K \in \mathcal{K}$ such that $L \sqsubseteq^{\emptyset} K$. We distinguish two cases:

- Assume that $K = D$. Since D is not subsumed by an atom in $\mathcal{L}$, it holds that $D \not\sqsubseteq^{\emptyset} L$.
- Now let $K \neq D$. Since $\mathcal{K}$ is a repair type, it does not contain $\sqsubseteq^{\emptyset}$-comparable atoms, which specifically implies that $D \not\sqsubseteq^{\emptyset} K$. Thus $D \not\sqsubseteq^{\emptyset} L$ must hold as otherwise D would be subsumed by K.

In both cases we conclude that $\mathsf{low}(\mathcal{K}, \mathcal{M}_D^0)$ contains either L itself or (if L is not maximal) an atom subsuming L, i.e., L is subsumed by an atom in $\mathsf{low}(\mathcal{K}, \mathcal{M}_D^0)$.

We proceed with the induction step ($i \to i+1$). Therefore let L be an atom in $\mathcal{L}$. Since $\mathcal{L}$ is a repair type for u, we have $\mathcal{B} \models L(u)$. Due to $\mathcal{L} \leq \mathcal{K}$ it further follows that L is subsumed by some atom K in $\mathcal{K}$. We show that $M \not\sqsubseteq^{\emptyset} L$ for each $M \in \mathcal{M}_D^{i+1}$. It then follows that $\mathsf{low}(\mathcal{K}, \mathcal{M}_D^{i+1})$ contains either L itself or an atom subsuming L, and thus L is subsumed by an atom in $\mathsf{low}(\mathcal{K}, \mathcal{M}_D^{i+1})$.

Assume to the contrary that there is an atom M in $\mathcal{M}_D^{i+1}$ such that $M \sqsubseteq^{\emptyset} L$. It cannot be the case that $M \in \mathcal{M}_D^i$ since this would lead to a contradiction with the induction hypothesis $\mathcal{L} \leq \mathsf{low}(\mathcal{K}, \mathcal{M}_D^i)$. Thus, consider the case where $M \in \mathcal{M}_D^{i+1} \setminus \mathcal{M}_D^i$. According to Definition 9 it follows that $M \in \mathsf{low}(\mathcal{K}, \mathcal{M}_D^i)$ and there is a subconcept $C \in \mathsf{Sub}(\mathcal{R}, \mathcal{T})$ with $\mathcal{B} \models C(u)$, $C \sqsubseteq^{\mathcal{T}} M$, and $\{C\} \not\leq \mathsf{low}(\mathcal{K}, \mathcal{M}_D^i)$. From $C \sqsubseteq^{\mathcal{T}} M$ and $M \sqsubseteq^{\emptyset} L$ it follows that $C \sqsubseteq^{\mathcal{T}} L$. Since $\mathcal{L}$ is a repair type for u, we infer that $\{C\} \leq \mathcal{L}$. Together with the induction hypothesis $\mathcal{L} \leq \mathsf{low}(\mathcal{K}, \mathcal{M}_D^i)$, this yields $\{C\} \leq \mathsf{low}(\mathcal{K}, \mathcal{M}_D^i)$, which is a contradiction.

Finally, recall that $\mathcal{M}_D$ is defined as $\mathcal{M}_D^j$ where j is the smallest index for which $\mathcal{M}_D^{j+1}$ equals $\mathcal{M}_D^j$. We thus obtain that $\mathcal{L} \leq \mathsf{low}(\mathcal{K}, \mathcal{M}_D)$. □

Using this lemma, we can now characterize non-minimality of an rsf.

Lemma 12. *A repair seed function on $\exists X.\mathcal{A}$ for $\mathcal{R}$ w.r.t. $\mathcal{T}$ is not $\leq$-minimal iff there exist an individual a and an atom $D \in s(a)$ such that $\{P\} \leq \mathsf{low}(s(a), \mathcal{M}_D)$ holds for each $P(a) \in \mathcal{R}$ with $\exists X.\mathcal{A} \models^{\mathcal{T}} P(a)$.*

Proof. If s is not $\leq$-minimal, then there is an rsf s' such that $s' < s$, i.e., there is $a \in \Sigma_\mathsf{I}$ such that $s'(a) < s(a)$. Since s' is an rsf, we have $\{P\} \leq s'(a)$ for all $P(a) \in \mathcal{R}$ with $\exists X.\mathcal{A} \models^{\mathcal{T}} P(a)$. By Lemma 11, $s'(a) < s(a)$ implies that there is $D \in s(a)$ such that $s'(a) \leq \mathsf{low}(s(a), \mathcal{M}_D)$. By transitivity, for each $P(a) \in \mathcal{R}$ with $\exists X.\mathcal{A} \models^{\mathcal{T}} P(a)$, we have $\{P\} \leq \mathsf{low}(s(a), \mathcal{M}_D)$.

To show the "if" direction, we construct a function $s' : \Sigma_\mathsf{I} \to \wp(\mathsf{Atoms}(\mathcal{R}, \mathcal{T}))$ such that $s'(b) := s(b)$ for each $b \in \Sigma_\mathsf{I} \setminus \{a\}$ and $s'(a) := \mathsf{low}(s(a), \mathcal{M}_D)$. By Lemma 10, $s'(a)$ is a repair type for a. Since for each $P(a) \in \mathcal{R}$ with $\exists X.\mathcal{A} \models^{\mathcal{T}} P(a)$, we have $\{P\} \leq \mathsf{low}(s(a), \mathcal{M}_D)$, we infer that s' is an rsf on $\exists X.\mathcal{A}$ for $\mathcal{R}$

w.r.t. $\mathcal{T}$. Since $s'(b) = s(b)$ for each $b \in \Sigma_\mathsf{I} \setminus \{a\}$ and $s'(a) < s(a)$, by Lemma 8, we infer that s is not $\leq$-minimal. □

Since there are linearly many atoms D in $s(a)$ and computing $\mathcal{M}_D$ and $\mathsf{low}(s(a), \mathcal{M}_D)$ can be done in polynomial time, we obtain the following complexity result.

Proposition 13. *$\leq$-minimality of repair seed functions is in* P.

Let us illustrate the decision procedure for non-minimality suggested by Lemma 12 by a small example.

Example 14. Consider the TBox $\mathcal{T} := \{\exists r.A_1 \sqsubseteq \exists r.A_2\}$, the quantified ABox $\exists X.\mathcal{A} := \exists\{x\}.\{r(a,x), A_1(x), A_2(x), B_1(x), B_2(x)\}$, and the repair request $\mathcal{R} := \{\exists r.(A_1 \sqcap B_1)(a), \exists r.(A_2 \sqcap B_2)(a)\}$. If we define a function $s : \Sigma_\mathsf{I} \to \wp(\mathsf{Atoms}(\mathcal{R}, \mathcal{T}))$ such that $s(a) = \{\exists r.A_1, \exists r.A_2\}$, then $s(a)$ is a repair type for a and s is a repair seed function on $\exists X.\mathcal{A}$ for $\mathcal{R}$ w.r.t. $\mathcal{T}$.

We use Lemma 12 to show that s is not $\leq$-minimal. For this purpose, we consider the atom $\exists r.A_1$ in $s(a)$, and construct the set $\mathcal{M}^0_{\exists r.A_1} := \{\exists r.A_1\}$. By Definition 7, we have $\mathsf{low}(s(a), \mathcal{M}^0_{\exists r.A_1}) = \{\exists r.(A_1 \sqcap B_1), \exists r.A_2\}$. However, this lowering set is not yet premise-saturated w.r.t. $\mathcal{T}$ since $\exists r.A_2$ is subsumed w.r.t. $\mathcal{T}$ by the subconcept $\exists r.A_1$, which is not subsumed w.r.t. $\emptyset$ by any atom from $\mathsf{low}(s(a), \mathcal{M}^0_{\exists r.A_1})$. By Definition 9, we thus add $\exists r.A_2$ to $\mathcal{M}^0_{\exists r.A_1}$, which yields the set $\mathcal{M}^1_{\exists r.A_1} := \{\exists r.A_1, \exists r.A_2\}$. The corresponding lowering set is $\mathsf{low}(s(a), \mathcal{M}^1_{\exists r.A_1}) = \{\exists r.(A_1 \sqcap B_1), \exists r.(A_2 \sqcap B_2)\}$. It is easy to see that this set is a repair repair type for a, which is strictly covered by $s(a)$. By looking at the repair request $\mathcal{R}$, we see that, for each concept assertion in $\mathcal{R}$, the respective concept is subsumed by some atom in $\mathsf{low}(s(a), \mathcal{M}^1_{\exists r.A_1})$. Thus, the condition on the right-hand side of the equivalence in Lemma 12 is satisfied for $\mathsf{low}(s(a), \mathcal{M}^1_{\exists r.A_1})$.

If we define $t(a) := \mathsf{low}(s(a), \mathcal{M}^1_{\exists r.A_1})$, then t is an rsf such that $t < s$. By Lemma 5, $\mathsf{rep}^{\mathcal{T}}_{\mathsf{IQ}}(\exists X.\mathcal{A}, t)$ strictly IQ-entails $\mathsf{rep}^{\mathcal{T}}_{\mathsf{IQ}}(\exists X.\mathcal{A}, s)$. For example, the former repair entails $(\exists r.A_1)(a)$ whereas the latter does not. This can be seen using Lemma 3.

4 Error-tolerant Reasoning w.r.t. Optimal Repairs

In error-tolerant reasoning, one does not commit to a single (classical or optimal) repair, but rather reasons w.r.t. all repairs. Brave entailment produces the consequences that are entailed by some repair whereas cautious entailment only produces consequences that are entailed by every repair. In the literature on inconsistency-tolerant and error-tolerant reasoning in the classical setting [10, 11, 15, 16, 18], IAR entailment (for "intersections of all repairs") is also considered, but in our setting of optimal repairs, where repairs are not necessarily subsets of the original ontology, it is not clear how to define this notion in an appropriate way.

If there is no repair, then everything is cautiously entailed and nothing is bravely entailed. We prevent this anomalous case by requiring that the repair request is solvable w.r.t. the given TBox.

Definition 15. *Let $\exists X.\mathcal{A}$ be a qABox, $\mathcal{T}$ an $\mathcal{EL}$ TBox, $\mathcal{R}$ a repair request that is solvable w.r.t. $\mathcal{T}$, and $\mathcal{Q}$ a finite set of $\mathcal{EL}$ concept assertions. Then $\mathcal{Q}$ is* bravely entailed *by $\exists X.\mathcal{A}$ w.r.t. $\mathcal{T}$ and $\mathcal{R}$ iff there is an optimal* IQ*-repair $\exists Z.\mathcal{C}$ of $\exists X.\mathcal{A}$ for $\mathcal{R}$ w.r.t. $\mathcal{T}$ such that $\exists Z.\mathcal{C} \models^{\mathcal{T}} C(a)$ for each $C(a) \in \mathcal{Q}$. It is* cautiously entailed *by $\exists X.\mathcal{A}$ w.r.t. $\mathcal{T}$ and $\mathcal{R}$ iff every optimal* IQ*-repair $\exists Z.\mathcal{C}$ of $\exists X.\mathcal{A}$ for $\mathcal{R}$ w.r.t. $\mathcal{T}$ satisfies $\exists Z.\mathcal{C} \models^{\mathcal{T}} C(a)$ for each $C(a) \in \mathcal{Q}$.*

In the following, we first show that brave entailment can be decided in polynomial time. For cautious entailment w.r.t. a TBox, the results proved in the previous section provide us with a coNP upper bound. Without a TBox, the complexity of cautious entailment drops to P.

4.1 Brave Entailment

The following lemma shows that brave entailment can be reduced to the instance problem in $\mathcal{EL}$.

Lemma 16. *The set of $\mathcal{EL}$ concept assertions $\mathcal{Q}$ is bravely entailed by $\exists X.\mathcal{A}$ for $\mathcal{R}$ w.r.t. $\mathcal{T}$ iff $\exists X.\mathcal{A} \models^{\mathcal{T}} \mathcal{Q}$ and no assertion in $\mathcal{P}$ is entailed by $\mathcal{Q}$ w.r.t. $\mathcal{T}$.*

Proof. If $\mathcal{Q}$ is bravely entailed, then there is an optimal IQ-repair $\exists Z.\mathcal{C}$ of $\exists X.\mathcal{A}$ for $\mathcal{P}$ w.r.t. $\mathcal{T}$ such that $\exists Z.\mathcal{C} \models^{\mathcal{T}} \mathcal{Q}$. Transitivity of entailment yields $\exists X.\mathcal{A} \models^{\mathcal{T}} \mathcal{Q}$. In addition, since $\exists Z.\mathcal{C}$ is a repair for $\mathcal{P}$, no assertion in $\mathcal{P}$ is entailed by $\exists Z.\mathcal{C}$ w.r.t. $\mathcal{T}$, and thus none can be entailed by $\mathcal{Q}$ w.r.t. $\mathcal{T}$.

Assume that $\exists X.\mathcal{A} \models^{\mathcal{T}} \mathcal{Q}$, and no assertion in $\mathcal{P}$ is entailed by $\mathcal{Q}$ w.r.t. $\mathcal{T}$. The set $\mathcal{Q}$ is an $\mathcal{EL}$ ABox, and thus there is a qABox $\exists Y.\mathcal{B}$ that is equivalent to $\mathcal{Q}$. Our assumptions on $\mathcal{Q}$ imply that $\exists Y.\mathcal{B}$ is an IQ-repair of $\exists X.\mathcal{A}$ for $\mathcal{R}$ w.r.t. $\mathcal{T}$. Since every repair is entailed by an optimal repair [3], there is an optimal IQ-repair $\exists Z.\mathcal{C}$ of $\exists X.\mathcal{A}$ for $\mathcal{P}$ w.r.t. $\mathcal{T}$ such that $\exists Z.\mathcal{C} \models^{\mathcal{T}} \exists Y.\mathcal{B}$, and thus $\exists Z.\mathcal{C} \models^{\mathcal{T}} \mathcal{Q}$. □

Since the instance problem in $\mathcal{EL}$ can be decided in polynomial time, this yields the following complexity result.

Theorem 17. *Brave entailment w.r.t. optimal* IQ*-repairs is in* P*.*

This approach for testing brave entailment can also be used to support computing a specific repair. In general, there may be exponentially many optimal repairs, but this set can be narrowed down by specifying not only consequences $\mathcal{R}$ to be removed, but also consequences $\mathcal{Q}$ that one wants to retain. Brave entailment can be used to check in polynomial time whether such a repair exists: in fact, Lemma 16 tells us that $\mathcal{Q}$ is bravely entailed by $\exists X.\mathcal{A}$ for $\mathcal{R}$ w.r.t. $\mathcal{T}$ iff the translation of $\mathcal{Q}$ into a qABox $\exists Y.\mathcal{B}$ is an IQ-repair of $\exists X.\mathcal{A}$ for $\mathcal{R}$ w.r.t. $\mathcal{T}$. In general, this repair will not be optimal. However, the next proposition shows that an rsf that induces an optimal repair entailing $\exists Y.\mathcal{B}$ (and thus also $\mathcal{Q}$) can be computed in polynomial time.

Proposition 18. *Let $\exists Y.\mathcal{B}$ be an IQ-repair of $\exists X.\mathcal{A}$ for $\mathcal{R}$ w.r.t. $\mathcal{T}$. Then we can compute in polynomial time a $\leq$-minimal rsf t such that $\mathsf{rep}_{\mathsf{IQ}}^{\mathcal{T}}(\exists X.\mathcal{A}, t) \models_{\mathsf{IQ}}^{\mathcal{T}} \exists Y.\mathcal{B}$. Since t is $\leq$-minimal, $\mathsf{rep}_{\mathsf{IQ}}^{\mathcal{T}}(\exists X.\mathcal{A}, t)$ is optimal.*

Proof. We know that every repair is entailed by a canonical repair. The proof of this fact (see proof of Proposition 8 in [4]) actually shows how to compute in polynomial time an rsf that induces this canonical repair. Thus, in the setting of our proposition, we can compute in polynomial time an rsf s such that $\mathsf{rep}_{\mathsf{IQ}}^{\mathcal{T}}(\exists X.\mathcal{A}, s) \models_{\mathsf{IQ}}^{\mathcal{T}} \exists Y.\mathcal{B}$. If s is $\leq$-minimal, then we are done. Otherwise, the proof of Lemma 12 tells us how to find an rsf s' such that $s' < s$. The rsf s' differs from s in the image for one individual a, where $s'(a) = \mathsf{low}(s(a), \mathcal{M}_D) < s(a)$ for an atom $D \in s(a)$. If s' is $\leq$-minimal, then we are done. Otherwise, we can compute an rsf s'' such that $s'' < s'$, etc. Since the next lemma implies that the length of such a chain $s > s' > s'' > \ldots$ is polynomially bounded by the number of individual names in $\exists X.\mathcal{A}$ and the cardinality of $\mathsf{Atoms}(\mathcal{R}, \mathcal{T})$, we reach a $\leq$-minimal rsf t with $t < s$ after a polynomial number of steps. By Lemma 5, $\mathsf{rep}_{\mathsf{IQ}}^{\mathcal{T}}(\exists X.\mathcal{A}, t) \models_{\mathsf{IQ}}^{\mathcal{T}} \mathsf{rep}_{\mathsf{IQ}}^{\mathcal{T}}(\exists X.\mathcal{A}, s)$, and thus $\mathsf{rep}_{\mathsf{IQ}}^{\mathcal{T}}(\exists X.\mathcal{A}, t) \models_{\mathsf{IQ}}^{\mathcal{T}} \exists Y.\mathcal{B}$. □

Lemma 19. *Let $\mathcal{S}$ be a set of $\mathcal{EL}$ concepts of cardinality m and $\mathcal{K}_0, \mathcal{K}_1, \ldots, \mathcal{K}_n$ be subsets of $\mathcal{S}$ such that $\mathcal{K}_0 > \mathcal{K}_1 > \ldots > \mathcal{K}_n$. Then $n \leq m$.*

Proof. For subsets $\mathcal{K}$ of $\mathcal{S}$, we define

$$\downarrow\mathcal{K} := \{\, C \mid C \in \mathcal{S} \text{ and } C \sqsubseteq^{\emptyset} K \text{ for some } K \in \mathcal{K} \,\}.$$

It is easy to see that $\mathcal{K} \leq \mathcal{L}$ iff $\downarrow\mathcal{K} \subseteq \downarrow\mathcal{L}$ holds for all subsets $\mathcal{K}, \mathcal{L}$ of $\mathcal{S}$. Thus $\mathcal{K}_0 > \mathcal{K}_1 > \ldots > \mathcal{K}_n$ implies $\downarrow\mathcal{K}_0 \supset \downarrow\mathcal{K}_1 \supset \ldots \supset \downarrow\mathcal{K}_n$. Since the cardinality of $\downarrow\mathcal{K}_0$ is bounded by the cardinality m of $\mathcal{S}$, this shows that $n \leq m$. □

Since, for solvable repair requests, the empty qABox $\exists\emptyset.\emptyset$ is a repair, Proposition 18 also yields the following result.

Corollary 20. *Let $\mathcal{T}$ be an $\mathcal{EL}$ TBox, $\exists X.\mathcal{A}$ a qABox, and $\mathcal{R}$ a repair request that is solvable w.r.t. $\mathcal{T}$. Then we can compute in polynomial time a $\leq$-minimal rsf t, which thus induces an optimal IQ-repair of $\exists X.\mathcal{A}$ for $\mathcal{R}$ w.r.t. $\mathcal{T}$.*

4.2 Cautious Entailment

Using the polynomiality results of Section 3, we can prove that cautious entailment is in coNP. For this, we show that non-entailment is in NP. To check whether $\mathcal{Q}$ is not cautiously entailed by $\exists X.\mathcal{A}$ w.r.t. $\mathcal{T}$ and $\mathcal{R}$, we guess a function $s : \Sigma_{\mathsf{I}} \to \wp(\mathsf{Atoms}(\mathcal{R}, \mathcal{T}))$ and check whether (i) s is a repair seed function; (ii) s is $\leq$-minimal; and (iii) there is $Q(a) \in \mathcal{Q}$ such that $\mathsf{rep}_{\mathsf{IQ}}^{\mathcal{T}}(\exists X.\mathcal{A}, s) \not\models^{\mathcal{T}} Q(a)$. Note that (i) can be decided in polynomial time by the definition of repair seed functions, (ii) by Proposition 13, and (iii) Proposition 4.

Theorem 21. *Cautious entailment w.r.t. optimal IQ-repairs is in* coNP.

Whether this upper bound is tight is still an open problem. If the TBox is empty, then we can show a polynomiality result.

The Case with an Empty TBox. We show the polynomial upper bound again for non-entailment, i.e., we try to find out whether there is an optimal repair that does not entail $\mathcal{Q}$. First note that, if $\mathcal{Q}$ is not entailed by $\exists X.\mathcal{A}$, then it cannot be entailed by an optimal repair. Thus, it is sufficient to concentrate on the case where $\exists X.\mathcal{A}$ entails $\mathcal{Q}$. For this case, the next lemma gives a characterization of non-entailment. While this characterization may look complicated, it is actually easy to see that its conditions can be checked in polynomial time. Intuitively, the reason why the case of the empty TBox is easier to handle is that then premise-saturatedness of repair types (see Definition 2) is trivially satisfied. More technically, this means that, in the characterization of non-minimality of a repair seed function in Lemma 12, the set $\mathcal{M}_D$ is equal to $\{D\}$, i.e., the iteration in Definition 9 terminates for $j = 0$. This gives us more control over how the sets $\mathsf{low}(s(a), \mathcal{M}_D)$ in Lemma 12 actually look like.

Lemma 22. *Let $\mathcal{Q}$ be a finite set of $\mathcal{EL}$ concept assertions such that $\exists X.\mathcal{A} \models \mathcal{Q}$. Then $\mathcal{Q}$ is not cautiously entailed by $\exists X.\mathcal{A}$ w.r.t. $\mathcal{R}$ iff there exist $C(a) \in \mathcal{Q}$, $D \in \mathsf{Atoms}(\mathcal{R})$, and $P(a) \in \mathcal{R}$ with $\mathcal{A} \models P(a)$ such that the following conditions are satisfied:*

1. *$P \sqsubseteq^\emptyset D$ and $C \sqsubseteq^\emptyset D$,*
2. *for each $D' \in \mathsf{Atoms}(\mathcal{R})$ with $D' \sqsubset^\emptyset D$ and $\mathcal{A} \models D'(a)$, we have $P \not\sqsubseteq^\emptyset D'$,*
3. *for each $P'(a) \in \mathcal{R} \setminus \{P(a)\}$ with $\mathcal{A} \models P'(a)$ and $P' \not\sqsubseteq^\emptyset D$, there is $E \in \mathsf{Atoms}(\mathcal{R})$ such that $P' \sqsubseteq^\emptyset E$ and $P \not\sqsubseteq^\emptyset E$.*

Proof. For the "only if" direction, if $\mathcal{Q}$ is not cautiously entailed by $\exists X.\mathcal{A}$ w.r.t. $\mathcal{R}$, then there exist $C(a) \in \mathcal{Q}$ and a $\leq$-minimal rsf s on $\exists X.\mathcal{A}$ for $\mathcal{R}$ such that $\mathsf{rep}_{\mathsf{IQ}}(\exists X.\mathcal{A}, s) \not\models C(a)$. By Lemma 3, the latter implies that there is $D \in s(a)$ such that $C \sqsubseteq^\emptyset D$.

Next, we show that there is $P(a) \in \mathcal{R}$ such that $P \sqsubseteq^\emptyset D$ and $\mathcal{A} \models P(a)$. Since s is $\leq$-minimal (for the case of an empty TBox), Lemma 12 implies that, for each $a \in \Sigma_\mathsf{I}$ and each $E \in s(a)$, there is $P_E(a) \in \mathcal{R}$ with $\mathcal{A} \models P_E(a)$ such that $\{P_E\} \not\leq \mathsf{low}(s(a), \{E\})$. However, $\{P_E\} \leq s(a)$ by the definition of repair seed functions. By Definition 7, the only atom from $s(a)$ that does not occur in $\mathsf{low}(s(a), \{E\})$ is E, which implies that $P_E \sqsubseteq^\emptyset E$. Consequently, there is $P(a) \in \mathcal{R}$ with $\mathcal{A} \models P(a)$ such that $P \sqsubseteq^\emptyset D$, which shows that Condition 1 of this lemma is satisfied by C, D, and P.

The construction of $\mathsf{low}(s(a), \{D\})$ removes D and replace it with those atoms $D' \in \mathsf{Atoms}(\mathcal{R})$ that are strictly subsumed by D such that $\mathcal{A} \models D'(a)$. However, $\{P\} \not\leq \mathsf{low}(s(a), \{D\})$ implies that, for each $D' \in \mathsf{Atoms}(\mathcal{R})$ with $D' \sqsubset^\emptyset D$ and $\mathcal{A} \models D'(a)$, we have $P \not\sqsubseteq^\emptyset D'$, i.e., Condition 2 is satisfied.

To show that Condition 3 is satisfied, we consider $P'(a) \in \mathcal{R} \setminus \{P(a)\}$ with $\mathcal{A} \models P'(a)$ and $P' \not\sqsubseteq^\emptyset D$. By the definition of an rsf, there must be $E \in s(a) \setminus \{D\}$ such that $P' \sqsubseteq^\emptyset E$. The fact that $\{P\} \not\leq \mathsf{low}(s(a), \{D\})$ implies that $P \not\sqsubseteq^\emptyset E$, which shows that Condition 3 of this lemma is indeed satisfied.

For the "if" direction, we assume that there exist $P(a) \in \mathcal{R}$ with $\mathcal{A} \models P(a)$ and $D \in \mathsf{Atoms}(\mathcal{R})$ such that all the three conditions of this lemma are satisfied.

We construct the set

$$\mathcal{K} := \{D\} \cup \mathsf{Max}_{\sqsubseteq^\emptyset}(\{E \in \mathsf{Atoms}(\mathcal{R}) \mid \text{there is } P'(a) \in (\mathcal{R} \setminus \{P(a)\}), \mathcal{A} \models P'(a), P' \not\sqsubseteq^\emptyset D, P' \sqsubseteq^\emptyset E, P \not\sqsubseteq^\emptyset E\}),$$

and show that it is a repair type. Since the TBox is empty, it suffices to consider only the first two properties of the definition of repair types (see Definition 2). The second property is immediately satisfied by the construction of $\mathcal{K}$. To show the first property, it is sufficient to prove that, for each $E \in \mathcal{K} \setminus \{D\}$, the atoms D and E are not $\sqsubseteq^\emptyset$-comparable. In fact, if $D \sqsubseteq^\emptyset E$, then $P \sqsubseteq^\emptyset E$, which contradicts our assumption that $E \in \mathcal{K} \setminus \{D\}$. If $E \sqsubseteq^\emptyset D$, then $P' \sqsubseteq^\emptyset D$ is a contradiction for some $P'(a) \in \mathcal{R} \setminus \{P(a)\}$, where $P' \sqsubseteq^\emptyset E$.

Using this set $\mathcal{K}$, we now define a function $s : \Sigma_\mathsf{I} \to \wp(\mathsf{Atoms}(\mathcal{R}))$ such that $s(a) := \mathcal{K}$ and $s(b) := \mathcal{M}_b$ for each individual $b \in \Sigma_\mathsf{I} \setminus \{a\}$, where $\mathcal{M}_b$ is a repair type for b and for each $R(b) \in \mathcal{R}$ with $\mathcal{A} \models R(b)$, there is $F \in \mathcal{M}_b$ such that $R \sqsubseteq^\emptyset F$. Such a repair type $\mathcal{M}_b$ exists for each $b \in \Sigma_\mathsf{I} \setminus \{a\}$ since $\mathcal{R}$ is solvable (see Proposition X in [4]).

We show that s is a repair seed function on $\exists X.\mathcal{A}$ for $\mathcal{R}$. For individuals $b \in \Sigma_\mathsf{I} \setminus \{a\}$, the condition on seed functions is satisfied, due to the way the sets $\mathcal{M}_b$ were chosen, i.e., such that $R(b) \in \mathcal{R}$ with $\mathcal{A} \models R(b)$ implies that there is an atom in $s(b)$ that subsumes R. We show that the corresponding condition also holds for $s(a)$. For $P(a)$, this is clear since is $D \in s(a)$ and $P \sqsubseteq^\emptyset D$. Furthermore, for each $P'(a) \in \mathcal{R} \setminus \{P(a)\}$, we distinguish two cases. If $P' \sqsubseteq^\emptyset D$, then we are done. Otherwise, by Condition 3, $P' \not\sqsubseteq^\emptyset D$ implies that there is $E \in \mathsf{Atoms}(\mathcal{R})$ such that $P' \sqsubseteq^\emptyset E$ and $P' \not\sqsubseteq^\emptyset E$. By the construction of $\mathcal{K}$, such an atom E occurs in $s(a) = \mathcal{K}$. This finally shows that s is an rsf on $\exists X.\mathcal{A}$ for $\mathcal{R}$.

Next, we show that, for each $\leq$-minimal rsf s' covered by s, the canonical repair induced by s' still does not entail $C(a)$. By Lemma 3, it is sufficient to show that $s'(a)$ contains an atom D' such that $D \sqsubseteq^\emptyset D'$. In fact, then $C \sqsubseteq^\emptyset D$ yields $C \sqsubseteq^\emptyset D'$ for $D' \in s'(a)$, and thus $C(a)$ is not entailed by the canonical repair induced by s', which is optimal since s' is minimal.

By contradiction, assume that there is a $\leq$-minimal rsf s' such that $s' \leq s$ and $D \not\sqsubseteq^\emptyset D'$ holds for all $D' \in s'(a)$. Thus, for each $D' \in s'(a)$, we have either $D' \sqsubset^\emptyset D$ or $D' \not\sqsubseteq^\emptyset D$. Consider again the concept P. Since s' is an rsf, there is $D' \in s'(a)$ such that $P \sqsubseteq^\emptyset D'$. Suppose that $D' \sqsubset^\emptyset D$. However, this is a contradiction since Condition 2 of this lemma states that P is not subsumed by any concept that is strictly subsumed by D. Otherwise, $D' \not\sqsubseteq^\emptyset D$. Since $s' \leq s$, we have $D' \sqsubseteq^\emptyset E$, where $E \in s(a) \setminus \{D\}$. By the definition of $\mathcal{K}$, P is not subsumed by E. However, $P \sqsubseteq^\emptyset D'$ and $D' \sqsubseteq^\emptyset E$, which yields a contradiction. □

This lemma reduces the non-entailment test to polynomially many subsumption and instance tests, each of which can be performed in polynomial time.

Theorem 23. *For an empty TBox, cautious entailment w.r.t. optimal* IQ-*repairs is in* P.

5 Conclusion

Inconsistency-tolerant and error-tolerant reasoning have been introduced in the DL literature [10, 11, 15, 16, 18] as a way to reason w.r.t. an inconsistent or erroneous ontology without having to commit to a specific repair. The usual entailment relations employed for this purpose are brave entailment (consequences entailed by some repair) and cautious entailment (consequences entailed by all repairs). In contrast to previous work, we use optimal repairs [3] instead of classical ones [9, 17, 19] when defining these relations. We investigated the complexity of the obtained entailment relations for the cases without and with a TBox, and could show a polynomial time upper bound for all cases except the one of cautious entailment with a TBox, for which we proved a coNP upper bound. The intuition underlying our use of optimal repairs is that a repair should not invent new consequences and should not have any of the unwanted consequences. A good repair should only remove consequences if this is required to achieve the other two goals.

Our approach for testing brave entailment can also be used to support computing a specific repair. In general, there may be exponentially many optimal repairs, but this set can be narrowed down by specifying not only consequences to be removed, but also ones that one wants to retain. We have shown that brave entailment can be used to check in polynomial time whether such a repair exists. In the positive case, we can compute in polynomial time a repair seed function that induces an optimal repair that entails all wanted consequences.

As pointed out in [16, 18], cautious entailment can be used to reason w.r.t. an erroneous ontology while waiting for a corrected update to be published by the organization that maintains this ontology. If the application is not repair but privacy preservation, one can use cautious entailment to define a censor [12] that prevents revealing certain secrets. The reason is that, in contrast to brave entailment, the set of cautious consequences is closed under (classical) entailment.

As future work, we will investigate whether our coNP upper bound for cautious entailment with a TBox is tight, and whether a notion of IAR entailment that is appropriate for optimal repairs can be found. We also intend to add support for role assertions both in the repair request and in the query. Furthermore, it would be interesting to consider error-tolerant reasoning w.r.t. the optimal TBox repairs in [13].

References

1. Baader, F., Brandt, S., Lutz, C.: Pushing the $\mathcal{EL}$ envelope. In: IJCAI-05, Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence. Professional Book Center (2005)
2. Baader, F., Horrocks, I., Lutz, C., Sattler, U.: An Introduction to Description Logic. Cambridge University Press (2017)
3. Baader, F., Koopmann, P., Kriegel, F., Nuradiansyah, A.: Computing optimal repairs of quantified ABoxes w.r.t. static $\mathcal{EL}$ TBoxes. In: Automated Deduction - CADE 28 - 28th International Conference on Automated Deduction, Proceedings. LNCS, vol. 12699. Springer (2021)

4. Baader, F., Koopmann, P., Kriegel, F., Nuradiansyah, A.: Computing optimal repairs of quantified ABoxes w.r.t. static $\mathcal{EL}$ TBoxes (extended version). LTCS-Report 21-01, Chair of Automata Theory, Institute of Theoretical Computer Science, Technische Universität Dresden, Dresden, Germany (2021)
5. Baader, F., Koopmann, P., Kriegel, F., Nuradiansyah, A.: Optimal ABox repair w.r.t. static $\mathcal{EL}$ TBoxes: From quantified ABoxes back to ABoxes. In: The Semantic Web - 19th International Conference, ESWC 2022, Proceedings. LNCS, vol. 13261, pp. 130–146. Springer (2022)
6. Baader, F., Koopmann, P., Kriegel, F., Nuradiansyah, A.: Optimal ABox repair w.r.t. static $\mathcal{EL}$ TBoxes: from quantified ABoxes back to ABoxes (extended version). LTCS-Report 22-01, Chair of Automata Theory, Institute of Theoretical Computer Science, Technische Universität Dresden, Dresden, Germany (2022)
7. Baader, F., Kriegel, F., Nuradiansyah, A.: Privacy-preserving ontology publishing for $\mathcal{EL}$ instance stores. In: Logics in Artificial Intelligence - 16th European Conference, JELIA 2019, Proceedings. LNCS, vol. 11468. Springer (2019)
8. Baader, F., Kriegel, F., Nuradiansyah, A., Peñaloza, R.: Computing compliant anonymisations of quantified ABoxes w.r.t. $\mathcal{EL}$ policies. In: The Semantic Web - ISWC 2020 - 19th International Semantic Web Conference, Proceedings, Part I. LNCS, vol. 12506. Springer (2020)
9. Baader, F., Suntisrivaraporn, B.: Debugging SNOMED CT using axiom pinpointing in the description logic $\mathcal{EL}^+$. In: Proceedings of the International Conference on Representing and Sharing Knowledge Using SNOMED (KR-MED'08). Phoenix, Arizona (2008)
10. Bienvenu, M., Rosati, R.: Tractable approximations of consistent query answering for robust ontology-based data access. In: IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence. IJCAI/AAAI (2013)
11. Calì, A., Lembo, D., Rosati, R.: On the decidability and complexity of query answering over inconsistent and incomplete databases. In: Proceedings of the Twenty-Second ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems. ACM (2003)
12. Cima, G., Lembo, D., Rosati, R., Savo, D.F.: Controlled query evaluation in description logics through instance indistinguishability. In: Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020. ijcai.org (2020)
13. Kriegel, F.: Optimal fixed-premise repairs of $\mathcal{EL}$ TBoxes. In: Proceedings of the 45th German Conference on Artificial Intelligence (KI 2022). LNCS, vol. 13404. Springer (2022)
14. Kriegel, F.: Optimal fixed-premise repairs of $\mathcal{EL}$ TBoxes (extended version). LTCS-Report 22-04, Chair of Automata Theory, Institute of Theoretical Computer Science, Technische Universität Dresden, Dresden, Germany (2022)
15. Lembo, D., Lenzerini, M., Rosati, R., Ruzzi, M., Savo, D.F.: Inconsistency-tolerant query answering in ontology-based data access. J. Web Semant. **33**, 3–29 (2015)
16. Ludwig, M., Peñaloza, R.: Error-tolerant reasoning in the description logic $\mathcal{EL}$. In: Logics in Artificial Intelligence - 14th European Conference, JELIA 2014. LNCS, vol. 8761. Springer (2014)
17. Parsia, B., Sirin, E., Kalyanpur, A.: Debugging OWL ontologies. In: Proc. of the 14th International Conference on World Wide Web (WWW'05). ACM (2005)
18. Peñaloza, R.: Error-tolerance and error management in lightweight description logics. Künstliche Intell. **34**(4), 491–500 (2020)
19. Schlobach, S., Huang, Z., Cornet, R., Harmelen, F.: Debugging incoherent terminologies. J. Automated Reasoning **39**(3), 317–349 (2007)