



Optimal ABox Repair w.r.t. Static \mathcal{EL} TBoxes: From Quantified ABoxes Back to ABoxes

Franz Baader, Patrick Koopmann, Francesco Kriegel,
and Adrian Nuradiansyah

Theoretical Computer Science, TU Dresden, Dresden, Germany
{franz.baader,patrick.koopmann,francesco.kriegel,
adrian.nuradiansyah}@tu-dresden.de

Abstract. Errors in Description Logic (DL) ontologies are often detected when a reasoner computes unwanted consequences. The question is then how to repair the ontology such that the unwanted consequences no longer follow, but as many of the other consequences as possible are preserved. The problem of computing such optimal repairs was addressed in our previous work in the setting where the data (expressed by an ABox) may contain errors, but the schema (expressed by an \mathcal{EL} TBox) is assumed to be correct. Actually, we consider a generalization of ABoxes called quantified ABoxes (qABoxes) both as input for and as result of the repair process. Using qABoxes for repair allows us to retain more information, but the disadvantage is that standard DL systems do not accept qABoxes as input. This raises the question, investigated in the present paper, whether and how one can obtain optimal repairs if one restricts the output of the repair process to being ABoxes. In general, such optimal ABox repairs need not exist. Our main contribution is that we show how to decide the existence of optimal ABox repairs in exponential time, and how to compute all such repairs in case they exist.

1 Introduction

Description Logics (DLs) [2] are a successful family of logic-based knowledge representation languages, which are employed in various application domains, but arguably their most prominent success was the adoption of the DL-based language OWL¹ as the standard ontology language for the Semantic Web. A DL knowledge base (aka ontology) consists of a TBox and an ABox. In the former, concepts can be used to state terminological constraints as so-called general concept inclusions (GCIs). For example, the concept $\exists\text{parent}.(Famous \sqcap Rich)$ describes individuals that have a parent that is both famous and rich, and the GCI $\exists\text{friend}.Famous \sqsubseteq Famous$ states that individuals that have a famous friend are famous themselves. The expressiveness of a DL depends on which constructors for building concepts are available. The concepts in our example use the

¹ <https://www.w3.org/OWL/>.

constructors conjunction (\sqcap) and existential restriction ($\exists r.C$), which together with the top concept (\top) are the ones available in the DL \mathcal{EL} , to which we restrict our attention here. While being quite inexpressive, \mathcal{EL} is nevertheless frequently used for building ontologies,² and it has the advantage over more expressive DLs that reasoning is polynomial w.r.t. \mathcal{EL} ontologies. In the ABox, one can relate named individuals with concepts and with each other. For example, the concept assertion ($\exists \text{parent.Rich}$)(*BEN*) states that Ben has a rich parent, and the role assertion *friend*(*BEN*, *JOHN*) says that Ben has John as friend. If concept assertions are restricted to employing only concept names, like *Famous*(*JOHN*), rather than complex concepts, then the ABox is called simple. DL systems provide their users with inference services that automatically derive implicit consequences such as instance relationships. For example, given the ABox assertions and the GCI introduced above, we can derive that Ben is famous, i.e., that the assertions *Famous*(*BEN*) follows from this ontology.

Although DL reasoners are usually sound (i.e., only derive instance relationship that indeed follow from the ontology), a computed consequence may still be incorrect in the application domain, due to the fact that the modelling of the domain in the ontology is erroneous. The question is then how to repair the ontology such that one gets rid of the unwanted consequences, but retains as many consequences as possible. Classical repair approaches that are based on removing axioms from the ontology [8, 11, 15, 16, 18, 19] are not optimal since, by removing large axioms, one may also lose information that does not contribute to the unwanted consequence. For example, if the concept assertion for John is (*Famous* \sqcap *Rich*)(*JOHN*) rather than just *Famous*(*JOHN*), then to get rid of the consequence *Famous*(*BEN*) we need to remove the whole assertion, and thus unnecessarily also lose the information that John is rich.

Extending on our previous work in [5, 7], we investigated in [3] how to compute optimal repairs in a setting where the ABox may contain errors, but the TBox is assumed to be a correct \mathcal{EL} TBox, and thus remains unchanged. More precisely, we consider a generalization of ABoxes called quantified ABoxes (qABoxes) both as input for and as result of the repair process since this allows us to retain more consequences. Such a qABox is a simple ABox where, however, some of the individuals are anonymized, which is formally expressed by existentially quantifying over them. In [3], we introduce two different notions of repair, depending on which entailment relation between qABoxes is considered: classical logical entailment or IQ-entailment, where the latter retains as many instance relationships as possible (but not necessarily answers to conjunctive queries). For the IQ case, we show that optimal IQ-repairs always exist and can be computed in exponential time. In the worst case, such repairs may be exponentially large and there may be exponentially many of them. Reusing an example from the introduction of [3], let us assume that the input ABox contains the information that Ben has a parent, Jerry, that is both rich and famous, that the TBox contains the GCI *Famous* \sqsubseteq *Rich*, and that we want to remove the consequence ($\exists \text{parent.}(Rich \sqcap Famous)$)(*BEN*). Using the optimized repair

² For example, the large medical ontology SNOMED CT is an \mathcal{EL} ontology.

approach of [3], we obtain the following qABox as one of the optimal IQ-repairs: $\exists\{y\}.\{parent(BEN, y), Rich(y), Famous(JERRY), Rich(JERRY)\}$.

The advantage of using qABoxes rather than ABoxes for repair is that more information can be retained (e.g. the fact that Ben has a rich parent). The disadvantage is that, though anonymized individuals are part of the OWL standard, DL systems usually do not accept them as input. Thus, the question arises whether one can also obtain optimal repairs if one restricts the output of the repair process to being ABoxes. In the above example, the qABox obtained as an optimal IQ-repair can actually be expressed by an ABox with complex concept assertions: $\{(\exists parent.Rich)(BEN), Famous(JERRY), Rich(JERRY)\}$.

However, this is not always the case. As an example, consider the ABox $\mathcal{A} := \{parent(BEN, JERRY), Rich(JERRY)\}$ and the TBox $\mathcal{T} := \{parent.Rich \sqsubseteq Famous, Famous \sqsubseteq \exists friend.Famous, \exists friend.Famous \sqsubseteq Famous\}$, which together imply that Ben is famous. Assume that Ben wants to get rid of this consequence. The repair approach of [3] yields the following qABox as an optimal IQ-repair: $\exists\{x, y\}.\{parent(BEN, x), Rich(JERRY), friend(BEN, y), friend(y, y)\}$. This qABox retains the information that Ben has a parent (but not that Jerry is this parent) and that Ben is the starting point of an infinite *friend*-chain, i.e., Ben belongs to the concepts $C_n := (\exists friend.)^n \top$ for all $n \geq 1$. The latter is the reason why this qABox cannot be expressed by an IQ-equivalent ABox, which in turn is the reason why there is no optimal ABox repair. The culprit is obviously the cycle *friend*(*y*, *y*). However, such cycles need not always cause problems. In fact, if we remove the third GCI $\exists friend.Famous \sqsubseteq Famous$ from the TBox, then the following qABox is an optimal IQ-repair:

$$\exists\{x, y\}.\{parent(BEN, x), Rich(JERRY), \\ friend(BEN, y), friend(y, y), Famous(y)\}.$$

This qABox can be expressed by an ABox that is IQ-equivalent to it w.r.t. the given TBox: $\{(\exists parent.\top)(BEN), Rich(JERRY), (\exists friend.Famous)(BEN)\}$. The reason is that, due to the existence of a famous friend of Ben, the GCI $Famous \sqsubseteq \exists friend.Famous$ now yields the infinite *friend*-chain.

These examples demonstrate that optimal ABox repairs may not always exist, and that it is not obvious to see when they do. The main contribution of the present paper is that we show how to decide the existence of optimal ABox repairs in exponential time, and how to compute all such repairs in case they exist. There may exist exponentially many such repairs, and each one may in the worst case be of double-exponential size. Our approach for showing these results roughly proceeds as follows. First, we observe that classical entailment between a qABox and an ABox coincides with so-called IRQ-entailment, which is slightly stronger than IQ-entailment by additionally taking role assertions between named individuals into account. Then, we show that both the canonical and the optimized IQ-repairs of [3] cannot only be used to obtain all optimal IQ-repairs, but also to compute all optimal IRQ-repairs. Subsequently, we introduce the notion of an optimal ABox approximation of a given qABox, and prove that the set of optimal ABox approximations of all optimal IRQ-repairs yields all

optimal ABox repairs. A given qABox may not have an optimal ABox approximation, but if it does, then this approximation is unique up to equivalence and of at most exponential size. Then we investigate the problem of deciding the existence of optimal ABox approximations. The first step is to transfer the qABox into a specific form, called pre-approximation, which is saturated w.r.t. the TBox and consists of the original role assertions between named individuals and for each named individual a a sub-qABox \mathcal{B}_a . We prove that the original qABox has an optimal ABox approximation iff all the named individuals a have a most specific concept C_a in \mathcal{B}_a w.r.t. the TBox. The optimal ABox approximation is then obtained by replacing each \mathcal{B}_a with $C_a(a)$ in the pre-approximation. We can then use the results stated in [20] to test the existence of the msc in polynomial time³ and to generate the at most exponentially large msc. Given that the optimal IRQ-repairs may be of exponential size, this yields the complexity upper bounds for testing the existence and computing optimal ABox repairs mentioned above. Due to space constraints, we cannot give complete proofs of all our results. They can be found in [4].

2 Preliminaries

We start with introducing the DL \mathcal{EL} as well as TBoxes and (quantified) ABoxes. Then we consider the entailment relations relevant for this paper.

The name space available for defining \mathcal{EL} concepts and ABox assertions is given by a *signature* Σ , which is the disjoint union of sets $\Sigma_{\mathcal{O}}$, $\Sigma_{\mathcal{C}}$, and $\Sigma_{\mathcal{R}}$ of *object names*, *concept names*, and *role names*. Starting with concept names and the top concept \top , \mathcal{EL} concepts are defined inductively: if C, D are \mathcal{EL} concepts and r is a role name, then $C \sqcap D$ (conjunction) and $\exists r.C$ (existential restriction) are also \mathcal{EL} concepts. An \mathcal{EL} general concept inclusion (GCI) is of the form $C \sqsubseteq D$, an \mathcal{EL} concept assertion is of the form $C(u)$, and a role assertion is of the form $r(u, v)$, where C, D are \mathcal{EL} concepts, $r \in \Sigma_{\mathcal{R}}$, and $u, v \in \Sigma_{\mathcal{O}}$. An \mathcal{EL} TBox is a finite set of \mathcal{EL} GCIs and an \mathcal{EL} ABox is a finite set of \mathcal{EL} concept assertions and role assertions. Such an ABox is called *simple* if all its concept assertions are of the form $A(u)$ with $A \in \Sigma_{\mathcal{C}}$. A *quantified ABox (qABox)* is of the form $\exists X.\mathcal{A}$ where X is a finite subset of $\Sigma_{\mathcal{O}}$ and \mathcal{A} is a simple ABox, which we call the *matrix* of $\exists X.\mathcal{A}$. We call the elements of X *variables* and the other object names occurring in \mathcal{A} *individuals*.⁴ The set of individual names occurring in $\exists X.\mathcal{A}$ is denoted with $\Sigma_I(\exists X.\mathcal{A})$, and the set of all object names (including the variables) with $\Sigma_{\mathcal{O}}(\exists X.\mathcal{A})$.

The semantics of the syntactic entities introduced above can either be defined directly using interpretations, or by a translation into first-order logic (FO). For the sake of brevity, we choose the latter approach (see [3] for the former). In the translation, the elements of $\Sigma_{\mathcal{O}}$, $\Sigma_{\mathcal{C}}$, and $\Sigma_{\mathcal{R}}$ are respectively viewed as constant

³ The proof for this polynomiality result in [20] is actually incorrect, but we show how to correct it.

⁴ The variables correspond to what we have called anonymized individuals in the introduction, and the individuals to what we have called named individuals.

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 $A \in \Sigma_C$ is translated into $A(x)$ and \top into $A(x) \vee \neg A(x)$ for an arbitrary $A \in \Sigma_C$;
- 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 $\exists y.(r(x, y) \wedge \phi_D(y))$, where $\phi_D(y)$ is obtained from $\phi_D(x)$ by replacing the free variable x by a different variable y .

GCI's $C \sqsubseteq D$ are translated into sentences $\phi_{C \sqsubseteq D} := \forall x.(\phi_C(x) \rightarrow \phi_D(x))$ and TBoxes \mathcal{T} into $\phi_{\mathcal{T}} := \bigwedge_{C \sqsubseteq D \in \mathcal{T}} \phi_{C \sqsubseteq D}$. Concept assertions $C(u)$ are translated into $\phi_C(u)$, role assertions $r(u, v)$ stay the same, and ABoxes \mathcal{A} are translated into the conjunction $\phi_{\mathcal{A}}$ of the translations of their assertions. For a quantified ABox $\exists X.\mathcal{A}$, the elements of X are viewed as first-order variables rather than constants, and its translation is $\exists \vec{x}.\phi_{\mathcal{A}}$, where \vec{x} is the tuple of the variables in X in arbitrary order.

Let α, β be (q)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 $(\phi_{\alpha} \wedge \phi_{\mathcal{T}}) \rightarrow \phi_{\beta}$ 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 . 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} . For ABoxes, the instance relation is defined analogously. Entailment between qABoxes w.r.t. an \mathcal{EL} TBox is NP-complete, but the subsumption and the instance problem are polynomial [7].

Note that ABoxes are a special case of qABoxes. For simple ABoxes, this is the case where $X = \emptyset$. For general ABoxes, one can express complex concept assertions by introducing existentially quantified variables (e.g., $\{(A \sqcap \exists r.B)(a)\}$ is equivalent to $\exists \{x\}.\{A(a), r(a, x), B(x)\}$). For this reason, the entailment relations defined below for qABoxes are also well-defined for ABoxes.

IQ-Entailment. If one is mainly interested in asking instance queries, i.e., in what kind of instance relations a qABox entails, then the following weaker form of entailment can be used [3, 7]. We say that the qABox $\exists X.\mathcal{A}$ *IQ-entails* the qABox $\exists Y.\mathcal{B}$ w.r.t. the \mathcal{EL} TBox \mathcal{T} (written $\exists X.\mathcal{A} \models_{\text{IQ}}^{\mathcal{T}} \exists Y.\mathcal{B}$) if every concept assertion $C(a)$ entailed w.r.t. \mathcal{T} by the latter is also entailed w.r.t. \mathcal{T} by the former. Whenever we compare two qABoxes $\exists X.\mathcal{A}$ and $\exists Y.\mathcal{B}$, we follow [7] and assume without loss of generality that they are *renamed apart*, which means that X is disjoint with $\Sigma_{\mathcal{O}}(\exists Y.\mathcal{B})$ and Y is disjoint with $\Sigma_{\mathcal{O}}(\exists X.\mathcal{A})$, and we further assume that the two qABoxes speak about the same set of individual names $\Sigma_1 := \Sigma_1(\exists X.\mathcal{A}) \cup \Sigma_1(\exists Y.\mathcal{B})$.

For the case of an empty TBox, it was shown in [7] that $\exists X.\mathcal{A} \models_{\text{IQ}}^{\emptyset} \exists Y.\mathcal{B}$ iff there is a simulation from $\exists Y.\mathcal{B}$ to $\exists X.\mathcal{A}$. A *simulation* from $\exists Y.\mathcal{B}$ to $\exists X.\mathcal{A}$ is a relation $\mathfrak{S} \subseteq \Sigma_{\mathcal{O}}(\exists Y.\mathcal{B}) \times \Sigma_{\mathcal{O}}(\exists X.\mathcal{A})$ such that $(a, a) \in \mathfrak{S}$ for each $a \in \Sigma_1$ and, for each $(u, v) \in \mathfrak{S}$, $A(u) \in \mathcal{B}$ implies $A(v) \in \mathcal{A}$ and $r(u, u') \in \mathcal{B}$

implies that there exists an object $v' \in \Sigma_{\mathcal{O}}(\exists X.\mathcal{A})$ such that $(u', v') \in \mathfrak{S}$ and $r(v, v') \in \mathcal{A}$. Since checking the existence of a simulation can be done in polynomial time [10], the simulation characterization of IQ-entailment shows that IQ-entailment between qABoxes can be decided in polynomial time if $\mathcal{T} = \emptyset$ [7].

- \sqcap -rule. If $(C_1 \sqcap \dots \sqcap C_n)(t) \in \mathcal{A}$, then remove this assertion from \mathcal{A} and add the assertions $C_1(t), \dots, C_n(t)$ to \mathcal{A} .
- \exists -rule. If $(\exists r.C)(t) \in \mathcal{A}$, then remove this assertion from \mathcal{A} , add the two assertions $r(t, x_C)$ and $C(x_C)$ to \mathcal{A} , and add x_C to X if it is not already there.
- \sqsubseteq -rule. If $t \in \Sigma_{\mathcal{O}}(\exists X.\mathcal{A})$, $C \sqsubseteq D \in \mathcal{T}$, $\mathcal{A} \models C(t)$, and $\mathcal{A} \not\models D(t)$, then add the assertion $D(t)$ to \mathcal{A} .

The \sqcap -rule has higher precedence than the \exists -rule, and the latter has higher precedence than the \sqsubseteq -rule.

Fig. 1. The IQ-saturation rules from [3].

To extend these results to the case of a non-empty TBox, the notion of an IQ-saturation is introduced in [3]. The saturation rules given in Fig. 1 add new variables and assertions to the qABox if the existence of a corresponding element and the validity of the assertion is implied by the TBox. To be more precise, for each existential restriction $\exists r.C$ occurring in \mathcal{T} , a fresh variable x_C not contained in the initial qABox is introduced. When applying the \exists -rule to an assertion of the form $(\exists r.C)(t)$, this variable is always used for the successor object. As pointed out in [3], IQ-saturation (i.e., the exhaustive application of the IQ-saturation rules) terminates in polynomial time and generates a qABox $\text{sat}_{\text{IQ}}^{\mathcal{T}}(\exists X.\mathcal{A})$, which can be seen as a qABox representation of what is called the *canonical model* in [13, Sect. 5.2]. IQ-entailment for qABoxes w.r.t. an \mathcal{EL} TBox is now characterized in [3] as follows.

Theorem 1 ([3]). *Let \mathcal{T} be an \mathcal{EL} TBox and $\exists X.\mathcal{A}$ and $\exists Y.\mathcal{B}$ qABoxes. Then the following statements are equivalent:*

- $\exists X.\mathcal{A} \models_{\text{IQ}}^{\mathcal{T}} \exists Y.\mathcal{B}$,
- $\text{sat}_{\text{IQ}}^{\mathcal{T}}(\exists X.\mathcal{A}) \models_{\text{IQ}}^{\emptyset} \exists Y.\mathcal{B}$,
- *there is a simulation from $\exists Y.\mathcal{B}$ to $\text{sat}_{\text{IQ}}^{\mathcal{T}}(\exists X.\mathcal{A})$.*

Since the IQ-saturation can be computed in polynomial time, this clearly shows that IQ-entailment for qABoxes w.r.t. an \mathcal{EL} TBox can also be decided in polynomial time.

IRQ-Entailment. If we are not only interested in implied concept assertions, but also in implied role assertions, then IQ-entailment is not sufficient. Instead, we must use IRQ-entailment. We say that the qABox $\exists X.\mathcal{A}$ *IRQ-entails* the qABox $\exists Y.\mathcal{B}$ w.r.t. the \mathcal{EL} TBox \mathcal{T} (written $\exists X.\mathcal{A} \models_{\text{IRQ}}^{\mathcal{T}} \exists Y.\mathcal{B}$) if every concept or role assertion entailed w.r.t. \mathcal{T} by the latter is also entailed w.r.t. \mathcal{T} by the former.

It is easy to see that a qABox cannot entail a role assertion involving a variable, and it can only entail a role assertion between individuals if its matrix contains this assertion. This yields the following characterization of IRQ-entailment, which shows that IRQ-entailment can be decided in polynomial time.

Proposition 2. *Let \mathcal{T} be an \mathcal{EL} TBox and $\exists X.\mathcal{A}$ and $\exists Y.\mathcal{B}$ qABoxes. Then the following statements are equivalent:*

- $\exists X.\mathcal{A} \models_{\text{IRQ}}^{\mathcal{T}} \exists Y.\mathcal{B}$,
- $\exists X.\mathcal{A} \models_{\text{IQ}}^{\mathcal{T}} \exists Y.\mathcal{B}$ and $r(a,b) \in \mathcal{B}$ implies $r(a,b) \in \mathcal{A}$ for all $r \in \Sigma_{\text{R}}$ and $a, b \in \Sigma_{\text{I}}$.

Since ABoxes consist of concept and role assertions, we obtain the following characterization of entailment between a qABox and an ABox, which implies that this entailment can be decided in polynomial time.

Proposition 3. *Let \mathcal{T} be an \mathcal{EL} TBox, $\exists X.\mathcal{A}$ a qABox, and \mathcal{B} an ABox. Then $\exists X.\mathcal{A} \models^{\mathcal{T}} \mathcal{B}$ iff $\exists X.\mathcal{A} \models_{\text{IRQ}}^{\mathcal{T}} \mathcal{B}$.*

3 Optimal ABox Repairs and Approximations

We first introduce the notion of an optimal repair w.r.t. an entailment relation, and show that the approaches for computing optimal IQ-repairs described in [3] can also be used to compute optimal IRQ-repairs. Then, we define optimal ABox approximations and show some useful properties for them. Finally, we introduce optimal ABox repairs, and describe how optimal ABox approximations can be used to obtain them from optimal IRQ-repairs.

3.1 Optimal IQ- and IRQ-Repairs

We start by recalling the definition of optimal repairs given in [3], but consider IRQ as an additional entailment relation.

Definition 4. *Let \mathcal{T} be an \mathcal{EL} TBox and $\text{QL} \in \{\text{IRQ}, \text{IQ}\}$.*

- An \mathcal{EL} repair request is a finite set of \mathcal{EL} concept assertions.
- Given a qABox $\exists X.\mathcal{A}$ and an \mathcal{EL} repair request \mathcal{R} , a QL-repair of $\exists X.\mathcal{A}$ for \mathcal{R} w.r.t. \mathcal{T} is a qABox $\exists Y.\mathcal{B}$ such that $\exists X.\mathcal{A} \models_{\text{QL}}^{\mathcal{T}} \exists Y.\mathcal{B}$ and $\exists Y.\mathcal{B} \not\models^{\mathcal{T}} C(a)$ for all $C(a) \in \mathcal{R}$.
- Such a repair $\exists Y.\mathcal{B}$ is optimal if there is no QL-repair $\exists Z.\mathcal{C}$ of $\exists X.\mathcal{A}$ for \mathcal{R} w.r.t. \mathcal{T} such that $\exists Z.\mathcal{C} \models_{\text{QL}}^{\mathcal{T}} \exists Y.\mathcal{B}$ and $\exists Y.\mathcal{B} \not\models_{\text{QL}}^{\mathcal{T}} \exists Z.\mathcal{C}$.

Two qABoxes are QL-equivalent if they QL-entail each other, and $\exists X.\mathcal{A}$ strictly QL-entails $\exists Y.\mathcal{B}$ if $\exists X.\mathcal{A} \models_{\text{QL}}^{\mathcal{T}} \exists Y.\mathcal{B}$ and $\exists Y.\mathcal{B} \not\models_{\text{QL}}^{\mathcal{T}} \exists X.\mathcal{A}$. We say that a set \mathfrak{R} of QL-repairs of $\exists X.\mathcal{A}$ for \mathcal{R} w.r.t. \mathcal{T} QL-covers all QL-repairs of $\exists X.\mathcal{A}$ for \mathcal{R} w.r.t. \mathcal{T} if for every QL-repair $\exists Y.\mathcal{B}$ of $\exists X.\mathcal{A}$ for \mathcal{R} w.r.t. \mathcal{T} there exists an element $\exists Z.\mathcal{C}$ of \mathfrak{R} such that $\exists Z.\mathcal{C} \models_{\text{QL}}^{\mathcal{T}} \exists Y.\mathcal{B}$. It is easy to see that such

a covering set \mathfrak{R} must contain, up to QL-equivalence, all optimal QL-repairs of $\exists X.\mathcal{A}$ for \mathcal{R} w.r.t. \mathcal{T} , and thus one can obtain from it, up to QL-equivalence, the set of all optimal QL-repairs of $\exists X.\mathcal{A}$ for \mathcal{R} w.r.t. \mathcal{T} by removing elements that are strictly QL-entailed by another element. Clearly, this set still QL-covers all QL-repairs of $\exists X.\mathcal{A}$ for \mathcal{R} w.r.t. \mathcal{T} .

In [3], two ways of computing such a covering set for IQ-repairs are described, the canonical IQ-repairs and the optimized IQ-repairs (see Proposition 8 and Theorem 14). Since these covering sets are of at most exponential cardinality, their elements are of at most exponential size, and IQ-entailment can be decided in polynomial time, this shows that, up to IQ-equivalence, the set of all optimal IQ-repairs can be computed in exponential time.

The canonical (optimized) IQ-repairs also yield covering sets for the IRQ case. The reason is basically that the approaches for constructing them introduced in [3] do not generate new role assertions between individuals and preserve as many of them as possible, although this is not required for IRQ-entailment.

Proposition 5. *Let \mathcal{T} be an \mathcal{EL} TBox, $\exists X.\mathcal{A}$ a qABox, and \mathcal{R} an \mathcal{EL} repair request. If \mathfrak{R} is the set of all canonical or all optimized IQ-repairs obtained from this input according to the definitions in [3], then \mathfrak{R} is a set of IRQ-repairs of $\exists X.\mathcal{A}$ for \mathcal{R} w.r.t. \mathcal{T} that IRQ-covers all IRQ-repairs of $\exists X.\mathcal{A}$ for \mathcal{R} w.r.t. \mathcal{T} . In particular, up to IRQ-equivalence, the set of optimal IRQ-repairs can be computed in exponential time, and it IRQ-covers all IRQ-repairs of $\exists X.\mathcal{A}$ for \mathcal{R} w.r.t. \mathcal{T} .*

Note that, though we have the same covering set \mathfrak{R} in the IQ and in the IRQ case, the sets of optimal repairs obtained from it by removing strictly entailed elements need not coincide since different entailment relations are used during this removal. Since the requirements for IQ entailment are weaker than for IRQ entailment, it could be that a qABox may be removed from \mathfrak{R} in the IQ case, but must be retained in IRQ case. Also notice that the proposition need not hold for arbitrary IQ-covering sets. Its proof uses properties of the canonical and the optimized IQ-repairs that need not hold for arbitrary covering sets.

Example 6. Consider the qABox $\exists\{x\}.\mathcal{A}$ for $\mathcal{A} = \{A(a), r(a, x), r(x, x)\}$, assume that the TBox is empty, and that the repair request is $\{A(a)\}$. An optimal IQ-repair $\exists\{x\}.\mathcal{A}'$ can be obtained from this qABox by removing the assertion $A(a)$ from \mathcal{A} , and this is also an optimal IRQ-repair. However, the ABox $\{r(a, a)\}$ is also an optimal IQ-repair since it is IQ-equivalent to $\exists\{x\}.\mathcal{A}'$, but it is not even an IRQ-repair since it is not IRQ-entailed by $\exists\{x\}.\mathcal{A}$.

3.2 Optimal ABox Approximations

Given a qABox we are now interested in finding an ABox that approximates it as closely as possible in the sense that a minimal amount of information is lost. In the definition below, we use classical entailment. But note that, according to Proposition 3, this coincides with IRQ-entailment.

Definition 7. *Given a qABox $\exists X.\mathcal{A}$ and an \mathcal{EL} TBox \mathcal{T} , we call an \mathcal{EL} ABox \mathcal{B} an ABox approximation of $\exists X.\mathcal{A}$ w.r.t. \mathcal{T} if $\exists X.\mathcal{A} \models^{\mathcal{T}} \mathcal{B}$. The ABox approximation \mathcal{B} of $\exists X.\mathcal{A}$ w.r.t. \mathcal{T} is optimal if there is no ABox approximation \mathcal{C} of $\exists X.\mathcal{A}$ w.r.t. \mathcal{T} such that $\mathcal{C} \models^{\mathcal{T}} \mathcal{B}$, but $\mathcal{B} \not\models^{\mathcal{T}} \mathcal{C}$.*

Such an optimal ABox approximation need not exist. The qABox $\exists\{x\}.\mathcal{A}'$ with $\mathcal{A}' = \{r(a, x), r(x, x)\}$ is an example for this case. In fact, this qABox entails $((\exists r.)^n \top)(a)$ for all $n \geq 1$, which is not possible for an ABox entailed by $\exists\{x\}.\mathcal{A}'$ since such an ABox cannot contain role assertions and can contain only finitely many concept assertions. However, if an optimal ABox approximation exists, then it is unique up to equivalence. This is an easy consequence of the fact that the union of two ABox approximations is again an ABox approximation.

Proposition 8. *If \mathcal{B}_1 and \mathcal{B}_2 are optimal ABox approximations of the qABox $\exists X.\mathcal{A}$ w.r.t. the \mathcal{EL} TBox \mathcal{T} , then \mathcal{B}_1 and \mathcal{B}_2 are equivalent w.r.t. \mathcal{T} .*

Optimal ABox approximations can now be characterized as follows.

Theorem 9. *The ABox \mathcal{B} is an optimal ABox approximation of $\exists X.\mathcal{A}$ w.r.t. \mathcal{T} iff $\exists X.\mathcal{A}$ and \mathcal{B} are IRQ-equivalent.*

Proof. First, assume that $\exists X.\mathcal{A}$ and \mathcal{B} are IRQ-equivalent w.r.t. \mathcal{T} . Then $\exists X.\mathcal{A} \models^{\mathcal{T}} \mathcal{B}$ by Proposition 3, and thus \mathcal{B} is an ABox approximation of $\exists X.\mathcal{A}$ w.r.t. \mathcal{T} . If \mathcal{C} is another ABox approximation of $\exists X.\mathcal{A}$ w.r.t. \mathcal{T} , then $\exists X.\mathcal{A} \models^{\mathcal{T}} \mathcal{C}$ by definition, and thus $\mathcal{B} \models^{\mathcal{T}} \mathcal{C}$ due to the assumed IRQ-equivalence. This shows optimality of \mathcal{B} .

Second, assume that \mathcal{B} is an optimal ABox approximation of $\exists X.\mathcal{A}$ w.r.t. \mathcal{T} that is not IRQ-equivalent with $\exists X.\mathcal{A}$. Then there is either a role assertion that belongs to \mathcal{A} , but not to \mathcal{B} , or a concept assertion that is entailed w.r.t. \mathcal{T} by $\exists X.\mathcal{A}$, but not by \mathcal{B} . Adding this assertion to \mathcal{B} yields an ABox \mathcal{B}' that is an ABox approximation of $\exists X.\mathcal{A}$ w.r.t. \mathcal{T} . In addition, it satisfies $\mathcal{B}' \models^{\mathcal{T}} \mathcal{B}$, but not $\mathcal{B} \models^{\mathcal{T}} \mathcal{B}'$, which contradicts the assumed optimality of \mathcal{B} . \square

An approach for deciding whether a given qABox has an optimal ABox approximation, and for computing it in case it exists, will be described in Sect. 4. But first, we show how optimal ABox approximations can be used to compute optimal ABox repairs.

3.3 Optimal ABox Repairs

The repair approaches developed in [3] in general yield quantified ABoxes as output, even if the input is an ABox. We are now interested in producing repairs that are ABoxes. The approach developed below does not require the input to be an ABox. It actually assumes that the input is a qABox, which means that input ABoxes first need to be transformed into equivalent qABoxes.

Definition 10. Let \mathcal{T} be an \mathcal{EL} TBox, $\exists X.\mathcal{A}$ a qABox, and \mathcal{R} an \mathcal{EL} repair request. We call an \mathcal{EL} ABox \mathcal{B} an ABox repair of $\exists X.\mathcal{A}$ for \mathcal{R} w.r.t. \mathcal{T} if $\exists X.\mathcal{A} \models^{\mathcal{T}} \mathcal{B}$ and $\mathcal{B} \not\models^{\mathcal{T}} C(a)$ for all $C(a) \in \mathcal{R}$. The ABox repair \mathcal{B} of $\exists X.\mathcal{A}$ for \mathcal{R} w.r.t. \mathcal{T} is optimal if there is no ABox repair \mathcal{C} of $\exists X.\mathcal{A}$ for \mathcal{R} w.r.t. \mathcal{T} such that $\mathcal{C} \models^{\mathcal{T}} \mathcal{B}$, but $\mathcal{B} \not\models^{\mathcal{T}} \mathcal{C}$.

Our approach for computing optimal ABox repairs proceeds as follows: first, we compute the set of all optimal IRQ-repairs of $\exists X.\mathcal{A}$, and then ABox-approximate the elements of this set. In the following, if we say that \mathfrak{R} is the set of optimal IRQ-repairs of a qABox, we mean that, for every optimal IRQ-repair, \mathfrak{R} contains one element of its IRQ-equivalence class. Also, for a given qABox $\exists Y.\mathcal{B}$, we define

$$\text{Oapp}^{\mathcal{T}}(\exists Y.\mathcal{B}) := \begin{cases} \{\mathcal{C}\} & \text{for an optimal ABox approx. } \mathcal{C} \text{ of } \exists Y.\mathcal{B} \text{ w.r.t. } \mathcal{T}, \\ \emptyset & \text{if no optimal ABox approx. of } \exists Y.\mathcal{B} \text{ w.r.t. } \mathcal{T} \text{ exists.} \end{cases}$$

Theorem 11. Let $\exists X.\mathcal{A}$ be a qABox, \mathcal{T} an \mathcal{EL} -TBox, \mathcal{R} an \mathcal{EL} repair request, and \mathfrak{R} the set of optimal IRQ-repairs of $\exists X.\mathcal{A}$ for \mathcal{R} w.r.t. \mathcal{T} . Then the set

$$\bigcup_{\exists Y.\mathcal{B} \in \mathfrak{R}} \text{Oapp}^{\mathcal{T}}(\exists Y.\mathcal{B})$$

consists of all optimal ABox repairs of $\exists X.\mathcal{A}$ for \mathcal{R} w.r.t. \mathcal{T} up to equivalence.

Proof. First, assume that the ABox \mathcal{C} belongs to the union defined in the statement of the theorem. Then $\exists X.\mathcal{A} \models_{\text{IRQ}}^{\mathcal{T}} \exists Y.\mathcal{B} \models^{\mathcal{T}} \mathcal{C}$ for some qABox $\exists Y.\mathcal{B} \in \mathfrak{R}$ that has \mathcal{C} as an optimal ABox approximation. This implies that \mathcal{C} does not entail any of the concept assertions in \mathcal{R} (since $\exists Y.\mathcal{B}$ does not) and that $\exists X.\mathcal{A} \models^{\mathcal{T}} \mathcal{C}$. Thus, \mathcal{C} is an ABox repair of $\exists X.\mathcal{A}$ for \mathcal{R} w.r.t. \mathcal{T} . It remains to show that it is optimal. Assume to the contrary that \mathcal{C}' is an ABox repair of $\exists X.\mathcal{A}$ for \mathcal{R} w.r.t. \mathcal{T} such that $\mathcal{C}' \models^{\mathcal{T}} \mathcal{C}$, but $\mathcal{C} \not\models^{\mathcal{T}} \mathcal{C}'$. Since \mathcal{C} and $\exists Y.\mathcal{B}$ are IRQ-equivalent by Theorem 9, this is a contradiction to the fact that $\exists Y.\mathcal{B}$ is an optimal IRQ-repair of $\exists X.\mathcal{A}$ for \mathcal{R} w.r.t. \mathcal{T} since \mathcal{C}' would then be a better IRQ-repair.

Second, assume that the ABox \mathcal{C} is an optimal ABox repair of $\exists X.\mathcal{A}$ for \mathcal{R} w.r.t. \mathcal{T} . Then \mathcal{C} is also an IRQ-repair of $\exists X.\mathcal{A}$ for \mathcal{R} w.r.t. \mathcal{T} , and thus Proposition 5 yields that there is an optimal IRQ-repair $\exists Y.\mathcal{B} \in \mathfrak{R}$ such that $\exists X.\mathcal{A} \models_{\text{IRQ}}^{\mathcal{T}} \exists Y.\mathcal{B} \models_{\text{IRQ}}^{\mathcal{T}} \mathcal{C}$. We know by Proposition 3 that the second IRQ-entailment is in fact an entailment, and thus \mathcal{C} is an ABox approximation of $\exists Y.\mathcal{B}$. It remains to show that it is optimal. Assume to the contrary that \mathcal{C}' is an ABox approximation of $\exists Y.\mathcal{B}$ such that $\exists Y.\mathcal{B} \models^{\mathcal{T}} \mathcal{C}' \models^{\mathcal{T}} \mathcal{C}$, but $\mathcal{C} \not\models^{\mathcal{T}} \mathcal{C}'$. But then \mathcal{C}' is an ABox repair of $\exists X.\mathcal{A}$ for \mathcal{R} w.r.t. \mathcal{T} (since $\exists Y.\mathcal{B}$ is a repair) that is better than \mathcal{C} , which contradicts our assumption that \mathcal{C} is optimal. \square

Once we have developed a method for computing the sets $\text{Oapp}^{\mathcal{T}}(\exists Y.\mathcal{B})$, this theorem shows how to compute the set of all optimal ABox repairs of a given qABox. Such a method will be introduced in the next section. Before doing this, we want to point out that, in contrast to the set of optimal IRQ-repairs, which covers all IRQ-repairs, the set of optimal ABox repairs in general does not cover all ABox repairs.

Example 12. Consider the ABox $\mathcal{A} = \{A(a), r(a, b), B(b)\}$, the TBox $\mathcal{T} = \{B \sqsubseteq \exists r. B, \exists r. B \sqsubseteq B\}$ and the repair request $\mathcal{R} = \{(A \sqcap \exists r. B)(a)\}$. There are basically three options for IRQ-repairing \mathcal{A} : remove $A(a)$, remove $B(b)$, or remove $r(a, b)$. Since things implied by the TBox must also be taken into account, these three options yield the following optimal IRQ-repairs of \mathcal{A} for \mathcal{R} w.r.t. \mathcal{T} :⁵ $\mathcal{B}_1 = \{r(a, b), B(b)\}$ as well as $\exists\{x\}.\mathcal{B}_i$ for $i = 2, 3$, where $\mathcal{B}_2 = \{A(a), r(a, b), r(b, x), r(x, x)\}$ and $\mathcal{B}_3 = \{A(a), B(b), r(a, x), r(x, x)\}$. Of these three, \mathcal{B}_1 is already an ABox, and thus its own optimal ABox approximation, whereas the other two have no optimal ABox approximation. However, they have non-optimal ABox approximations, which are not necessarily covered by \mathcal{B}_1 . For example, $\{A(a), r(a, b), (\exists r. \exists r. \top)(b)\}$ is an ABox approximation of $\exists\{x\}.\mathcal{B}_2$ and an ABox repair of \mathcal{A} for \mathcal{R} w.r.t. \mathcal{T} , but since it contains $A(a)$, it is not entailed by \mathcal{B}_1 .

4 Computing Optimal ABox Approximations

In this section, we assume that $\exists X.\mathcal{A}$ is a qABox and \mathcal{T} an \mathcal{EL} TBox. We will develop an approach for deciding whether $\exists X.\mathcal{A}$ has an optimal ABox approximation w.r.t. \mathcal{T} , which in the affirmative case also yields such an optimal approximation.

The first step is to saturate $\exists X.\mathcal{A}$ using the IQ-saturation rules of Fig. 1. In the following, let $\text{sat}_{\text{IQ}}^{\mathcal{T}}(\exists X.\mathcal{A})$ denote a (fixed) qABox obtained by applying the IQ-saturation rules exhaustively to $\exists X.\mathcal{A}$. Note that the size of $\text{sat}_{\text{IQ}}^{\mathcal{T}}(\exists X.\mathcal{A})$ is polynomial in the size of the input $\exists X.\mathcal{A}$ and \mathcal{T} , and that $\exists X.\mathcal{A}$ and $\text{sat}_{\text{IQ}}^{\mathcal{T}}(\exists X.\mathcal{A})$ are IQ-equivalent w.r.t. \mathcal{T} by Theorem 1. In addition, it is easy to see that these two qABoxes contain the same individuals and the same role assertions between individuals. Thus, they are even IRQ-equivalent w.r.t. \mathcal{T} . As before, we use Σ_1 to denote set of individuals of $\exists X.\mathcal{A}$.

In the next step, we transform $\text{sat}_{\text{IQ}}^{\mathcal{T}}(\exists X.\mathcal{A})$ into a new qABox, called pre-approximation, whose matrix basically consists of the union of ABoxes \mathcal{B}_a for each $a \in \Sigma_1$, extended with the role assertions between individuals in \mathcal{A} . Each ABox \mathcal{B}_a contains a as the only individual name, and further contains a fully anonymized copy of the saturation $\text{sat}_{\text{IQ}}^{\mathcal{T}}(\exists X.\mathcal{A})$, which is connected with a by indispensable role assertions.

Definition 13. *We call a role assertion $r(a, u)$ in $\text{sat}_{\text{IQ}}^{\mathcal{T}}(\exists X.\mathcal{A})$ for $a \in \Sigma_1$ indispensable if there is no role assertion $r(a, b)$ for $b \in \Sigma_1$ such that there is a simulation from $\text{sat}_{\text{IQ}}^{\mathcal{T}}(\exists X.\mathcal{A})$ to itself that contains (u, b) .*

Since an individual always simulates itself, only role assertion $r(a, u)$ where u is a variable can be indispensable. We are now ready to define the pre-approximation.

⁵ The IQ-repairs computed by the approaches in [3] would contain more assertions, which are however redundant for IRQ-entailment w.r.t. \mathcal{T} .

Definition 14. The pre-approximation $\text{pre-approx}_{\text{IRQ}}^{\mathcal{T}}(\exists X.\mathcal{A})$ of $\exists X.\mathcal{A}$ w.r.t. \mathcal{T} is defined as the quantified ABox $\exists Y.\mathcal{B}$, where

$$\begin{aligned} Y &:= \{u' \mid u \text{ is an object name occurring in } \underline{\text{sat}}_{\text{IRQ}}^{\mathcal{T}}(\exists X.\mathcal{A})\}, \\ \mathcal{B} &:= \bigcup \{ \mathcal{B}_a \mid a \text{ is an individual name in } \Sigma_1 \} \\ &\quad \cup \{ r(a, b) \mid r(a, b) \text{ occurs in } \underline{\text{sat}}_{\text{IRQ}}^{\mathcal{T}}(\exists X.\mathcal{A}) \text{ where } a, b \in \Sigma_1 \}, \\ \mathcal{B}_a &:= \{ A(a) \mid A(a) \text{ occurs in } \underline{\text{sat}}_{\text{IRQ}}^{\mathcal{T}}(\exists X.\mathcal{A}) \} \\ &\quad \cup \{ r(a, u') \mid r(a, u) \text{ occurs in } \underline{\text{sat}}_{\text{IRQ}}^{\mathcal{T}}(\exists X.\mathcal{A}) \text{ and is indispensable} \} \\ &\quad \cup \{ A(u') \mid A(u) \text{ occurs in } \underline{\text{sat}}_{\text{IRQ}}^{\mathcal{T}}(\exists X.\mathcal{A}) \} \\ &\quad \cup \{ r(u', v') \mid r(u, v) \text{ occurs in } \underline{\text{sat}}_{\text{IRQ}}^{\mathcal{T}}(\exists X.\mathcal{A}) \}. \end{aligned}$$

Obviously, the pre-approximation can be computed in polynomial time. In addition, it is IRQ-equivalent to $\underline{\text{sat}}_{\text{IRQ}}^{\mathcal{T}}(\exists X.\mathcal{A})$ [4].

Lemma 15. The q ABoxes $\underline{\text{sat}}_{\text{IRQ}}^{\mathcal{T}}(\exists X.\mathcal{A})$ and $\text{pre-approx}_{\text{IRQ}}^{\mathcal{T}}(\exists X.\mathcal{A})$ are IRQ-equivalent w.r.t. the empty TBox \emptyset , and thus also w.r.t. \mathcal{T} .

Since we already know that $\exists X.\mathcal{A}$ and $\underline{\text{sat}}_{\text{IRQ}}^{\mathcal{T}}(\exists X.\mathcal{A})$ are IRQ-equivalent w.r.t. \mathcal{T} , this shows that $\exists X.\mathcal{A}$ is IRQ-equivalent to its pre-approximation w.r.t. \mathcal{T} . Consequently, an ABox \mathcal{C} is an optimal ABox approximation of $\exists X.\mathcal{A}$ w.r.t. \mathcal{T} iff it is one of the pre-approximation w.r.t. \mathcal{T} .

To test whether $\text{pre-approx}_{\text{IRQ}}^{\mathcal{T}}(\exists X.\mathcal{A})$ has an optimal ABox approximation w.r.t. \mathcal{T} , it is sufficient to check whether, for all $a \in \Sigma_1$, the individual a has a most specific concept in \mathcal{B}_a w.r.t. \mathcal{T} .

Definition 16. Let \mathcal{C} be an \mathcal{EL} ABox, \mathcal{T} an \mathcal{EL} TBox, and a an individual name. The \mathcal{EL} concept C is a most specific concept (msc) of a in \mathcal{C} w.r.t. \mathcal{T} if $\mathcal{C} \models^{\mathcal{T}} C(a)$ and $\mathcal{C} \models^{\mathcal{T}} D(a)$ implies $C \sqsubseteq^{\mathcal{T}} D$ for all \mathcal{EL} concepts D .

The most specific concept need not exist, but if it does, then it is unique up to equivalence w.r.t. \mathcal{T} . The ABox $\mathcal{C} := \{r(a, a)\}$ is a simple example where the msc of a does not exist w.r.t. the empty TBox. In fact, $\mathcal{C} \models (\exists r.)^n \top(a)$ for all $n \geq 1$, and it is easy to see that no \mathcal{EL} concept can be subsumed by these infinitely many concepts. Note, however, that \mathcal{C} has an optimal ABox approximation since it is itself an ABox. In this case, the pre-approximation is $\{r(a, a)\} \cup \mathcal{B}_a$ where $\mathcal{B}_a = \{r(a', a')\}$. There is no role assertion $r(a, a')$ since $r(a, a)$ is not indispensable. While a' does not have an msc in \mathcal{B}_a , this is not what we are interested in. We want to know whether a has one, and the answer is “yes” since \top is an msc of a in \mathcal{B}_a . The problem of testing for the existence of and computing the msc in \mathcal{EL} was investigated in [20], where the following result is stated.

Proposition 17 ([20]). Let \mathcal{C} be an \mathcal{EL} ABox, \mathcal{T} an \mathcal{EL} TBox, and a an individual name. It can be decided in polynomial time whether a has a most specific concept in \mathcal{C} w.r.t. \mathcal{T} , and if the msc exists, then it can be computed in exponential time.

The main idea underlying the proof of this proposition (rephrased into the setting of the present paper) is to *unravel* the IQ-saturation of \mathcal{C} w.r.t. \mathcal{T} into a concept C_k an increasing number k of steps, starting from a . After each step, one tests whether the ABox $\{C_k(a)\}$ IQ-entails $\exists X.\mathcal{C}$ w.r.t. \mathcal{T} , where X consists of the object names in \mathcal{C} different from a . In case this test succeeds, the concept C_k is the msc of a in \mathcal{C} w.r.t. \mathcal{T} . This yields an effective test for the existence of the msc since the following can be shown: there is a polynomial p such that the entailment test succeeds after at most $p(|\mathcal{C}|, |\mathcal{T}|)$ steps iff the msc exists.

For example, for the ABox $\mathcal{C}^{(1)} = \{r(a, a)\}$ and the TBox $\mathcal{T}^{(1)} = \emptyset$, the 0-step unraveling is $C_0^{(1)} = \top$, the 1-step unraveling is $C_1^{(1)} = \exists r.\top$, the two-step unraveling is $C_2^{(1)} = \exists r.\exists r.\top$, etc. It is easy to see that there is no k such that the entailment test succeeds. Thus, it does not succeed for $k(\mathcal{C}^{(1)}, \mathcal{T}^{(1)})$, which shows that a does not have an msc. If instead we consider the ABox $\mathcal{C}^{(2)} = \{A(a), r(a, b), s(a, b), r(b, c), s(b, c), B(c)\}$ w.r.t. $\mathcal{T}^{(2)} = \emptyset$, then the 0-step unraveling is $C_0^{(2)} = A$, the 1-step unraveling is $C_1^{(2)} = A \sqcap \exists r.\top \sqcap \exists s.\top$, the 2-step unraveling is $C_2^{(2)} = A \sqcap \exists r.(\exists r.B \sqcap \exists s.B) \sqcap \exists s.(\exists r.B \sqcap \exists s.B)$, and the 3-step unraveling is identical to $C_2^{(2)}$. The entailment test succeeds for $k = 2$. It is easy to see that, whenever the unraveling becomes stable (which happens if no cycle in the ABox is reachable from a), then the entailment test succeeds. However, a reachable cycle in the ABox need not prevent the existence of the msc. For example, the individual a has the msc $\exists r.B$ in $\mathcal{C}^{(3)} = \{r(a, b), r(b, b), B(b)\}$ w.r.t. $\mathcal{T}^{(3)} = \{B \sqsubseteq \exists r.B\}$.

As sketched until now, this method for deciding the existence of the msc does not yield a polynomial-time decision procedure. The reason is that, though the bound $k(\mathcal{C}, \mathcal{T})$ on the number of steps is polynomial, the unraveled concepts C_k may become exponential even for $k \leq k(\mathcal{C}, \mathcal{T})$, as can be seen using an obvious generalization of our example ABox $\mathcal{C}^{(2)}$. This problem can be avoided by employing structure-sharing, which can be realized by representing the ABoxes $\{C_k(a)\}$ by IQ-equivalent qABoxes. In our second example, the ABox $\{C_2^{(2)}(a)\}$ can be represented by the more compact IQ-equivalent qABox $\exists \{x, y\}. \{A(a), r(a, x), s(a, x), r(x, y), s(x, y), B(y)\}$ (see the definition of the k -unraveling in [1] for how such an unraveling with structure sharing can be defined in general). It is easy to see that the qABoxes representing the ABoxes $\{C_k(a)\}$ are of polynomial size. Since IQ-entailment between qABoxes is polynomial, this yields the polynomiality result stated in the proposition. Note, however, that the msc obtained this way is still an unraveled concept C_k without structure sharing, and thus may be of exponential size.

The following theorem shows that existence of the optimal ABox approximation can be reduced to existence of the msc (see [4] for the proof).

Theorem 18. *Let $\exists X.\mathcal{A}$ be a qABox with set of individuals Σ_1 , let \mathcal{T} be an \mathcal{EL} TBox, and let \mathcal{B}_a for all $a \in \Sigma_1$ be the ABoxes introduced in Definition 14. Then $\exists X.\mathcal{A}$ has an optimal ABox approximation w.r.t. \mathcal{T} iff, for all individuals*

$a \in \Sigma_1$, the msc of a in \mathcal{B}_a w.r.t. \mathcal{T} exists. If the latter condition is satisfied and C_a are these most specific concepts, then the following ABox is an optimal ABox approximation of $\exists X.\mathcal{A}$ w.r.t. \mathcal{T} :

$$\{C_a(a) \mid a \in \Sigma_1\} \cup \{r(a,b) \mid r(a,b) \text{ occurs in } \underline{\text{sat}}_{\text{IQ}}^{\mathcal{T}}(\exists X.\mathcal{A}) \text{ where } a,b \in \Sigma_1\}.$$

In particular, the existence of an optimal ABox approximation can be tested in polynomial time and such an optimal approximation can be computed in exponential time if it exists.

5 Computing Optimal ABox Repairs

We can now reap the benefits from the results shown in the previous two sections. Given a qABox $\exists X.\mathcal{A}$, an \mathcal{EL} TBox \mathcal{T} , and an \mathcal{EL} repair request \mathcal{R} , we can compute the set \mathfrak{R} of optimal IRQ-repairs of $\exists X.\mathcal{A}$ for \mathcal{R} w.r.t. \mathcal{T} in exponential time. More precisely, by Proposition 5 this set contains at most exponentially many repairs, each of which has at most exponential size. Theorem 11 then says that the set of all optimal ABox repairs of $\exists X.\mathcal{A}$ for \mathcal{R} w.r.t. \mathcal{T} (up to equivalence) consists of the optimal ABox approximations w.r.t. \mathcal{T} of those elements of \mathfrak{R} for which such an optimal approximation exists. Finally, Theorem 18 shows how to decide existence of such optimal approximations and how to compute them if they exist. Since the elements of \mathfrak{R} are already of exponential size, existence can be tested in exponential time and the size of the computed approximations is at most double-exponential.

Theorem 19. *Let $\exists X.\mathcal{A}$ be a qABox, \mathcal{T} an \mathcal{EL} -TBox, and \mathcal{R} an \mathcal{EL} repair request. Then the existence of an optimal ABox repair of $\exists X.\mathcal{A}$ for \mathcal{R} w.r.t. \mathcal{T} can be decided in exponential time, and the set of all such repairs can be computed in double-exponential time. This set contains at most exponentially many elements, each of which has at most double-exponential size.*

If the given qABox does not have an optimal repair or if we are looking for a repair not covered by an optimal one, our approach can also be used to compute non-optimal ABox repairs. In fact, consider an optimal IRQ-repair that does not have an optimal ABox approximation. Then there are individuals a whose msc in \mathcal{B}_a does not exist. Following [17], we can then use the role-depth bounded msc instead, which is basically obtained by unraveling up to a fixed bound k on the role-depth (i.e., the maximal nesting of existential restrictions). This way, we can produce a set of (possibly) non-optimal ABox repairs, which covers all ABox repairs whose concept assertions satisfy this bound on the role depth.

There are also cases where the existence of the optimal ABox approximation of the optimal IRQ-repairs is guaranteed. In fact, if the qABox is acyclic and the TBox is cycle-restricted (i.e., there is no concept C such that $C \sqsubseteq^{\mathcal{T}} \exists r_1. \dots \exists r_k. C$, as defined in [3]), then the optimal IRQ-repairs are acyclic, which implies that the ABoxes \mathcal{B}_a in the pre-approximations are also acyclic. Consequently, all optimal IRQ-repairs have an optimal ABox approximation. The following corollary is an easy consequence of this observation.

Corollary 20. *Let $\exists X.A$ be an acyclic qABox, \mathcal{T} a cycle-restricted \mathcal{EL} -TBox, and \mathcal{R} an \mathcal{EL} repair request. Then the set of optimal ABox repairs of $\exists X.A$ for \mathcal{R} w.r.t. \mathcal{T} is non-empty, and it IRQ-covers all ABox repairs of $\exists X.A$ for \mathcal{R} w.r.t. \mathcal{T} .*

6 Conclusion

Traditional repair approaches for DL-based ontologies, which compute maximal subsets of the ontology that do not have the unwanted consequences, are syntax-dependent and thus may remove too many consequences. Recently developed syntax-independent approaches for repairing DL ABoxes [3, 5, 7] compute optimal repairs that do not lose consequences unnecessarily, but they have the disadvantage that they produce quantified ABoxes rather than traditional ABoxes. In this paper we show how to overcome this problem by developing methods for computing optimal repairs that are traditional ABoxes. These methods are based on the computation of optimal IRQ-repairs, by adapting the approaches in [3] for computing optimal IQ-repairs, and then optimally approximating these qABoxes with ABoxes.

A perceived disadvantage of our approach could be that optimal ABox repairs need not exist, and even if they do, they need not cover all ABox repairs. However, by Corollary 20 this problem does not occur if the ABox is acyclic and the TBox is cycle-restricted. To see how often this corollary applies in practice, we checked the 80 large ontologies used in the experiments in [3]: 62 have cycle-restricted TBoxes, and of those only 7 have cyclic ABoxes. Thus, our Corollary 20 applies to 55 of the 80 ontologies considered in [3].

Another disadvantage could be the potentially double-exponential size of optimal ABox repairs. However, the first exponential comes from the computation of the optimal IQ-repairs, and the experiments in [3] indicate that this exponential blow-up does not occur in practice if the optimized approach for computing IQ-repairs is used. We do not yet have experimental results regarding the possible exponential blow-up due to the computation of ABox approximations, but would be surprised if this happened often in practice.

What is called “repair” in the DL community is closely related to what is called “contraction” in the Belief Change community. For classical repairs and also for the gentle repairs of [6], this connection was investigated in [14]. It would be interesting to see whether this investigation can be extended to our optimal ABox repairs. The original intention underlying our repair approach is that the ontology engineer chooses one of the computed optimal repairs as the new, repaired ABox. Alternatively, one could try to adapt the different repair semantics employed in inconsistency-tolerant query answering [9, 12] from classical repairs to our optimal repairs.

Acknowledgements. This work was partially supported by the AI competence center ScaDS.AI Dresden/Leipzig and the Deutsche Forschungsgemeinschaft (DFG), Grant 430150274, and Grant 389792660 within TRR 248.

References

1. Baader, F.: A graph-theoretic generalization of the least common subsumer and the most specific concept in the description logic \mathcal{EL} . In: Hromkovič, J., Nagl, M., Westfechtel, B. (eds.) WG 2004. LNCS, vol. 3353, pp. 177–188. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-30559-0_15
2. Baader, F., Horrocks, I., Lutz, C., Sattler, U.: An Introduction to Description Logic. Cambridge University Press, Cambridge (2017). <https://doi.org/10.1017/9781139025355>
3. Baader, F., Koopmann, P., Kriegel, F., Nuradiansyah, A.: Computing optimal repairs of quantified ABoxes w.r.t. static \mathcal{EL} TBoxes. In: Platzer, A., Sutcliffe, G. (eds.) CADE 2021. LNCS (LNAI), vol. 12699, pp. 309–326. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-79876-5_18
4. 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). <https://doi.org/10.25368/2022.65>
5. Baader, F., Kriegel, F., Nuradiansyah, A.: Privacy-preserving ontology publishing for \mathcal{EL} instance stores. In: Calimeri, F., Leone, N., Manna, M. (eds.) JELIA 2019. LNCS (LNAI), vol. 11468, pp. 323–338. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-19570-0_21
6. Baader, F., Kriegel, F., Nuradiansyah, A., Peñaloza, R.: Making repairs in description logics more gentle. In: Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018, Tempe, Arizona, 30 October–2 November 2018, pp. 319–328. AAAI Press (2018). <https://aaai.org/ocs/index.php/KR/KR18/paper/view/18056>
7. Baader, F., Kriegel, F., Nuradiansyah, A., Peñaloza, R.: Computing compliant anonymisations of quantified ABoxes w.r.t. \mathcal{EL} policies. In: Pan, J.Z., et al. (eds.) ISWC 2020. LNCS, vol. 12506, pp. 3–20. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-62419-4_1
8. Baader, F., Suntisrivaraporn, B.: Debugging SNOMED CT using axiom pinpointing in the description logic \mathcal{EL}^+ . In: Proceedings of the Third International Conference on Knowledge Representation in Medicine, Phoenix, Arizona, USA, 31 May–2 June 2008. CEUR Workshop Proceedings, vol. 410. CEUR-WS.org (2008). <http://ceur-ws.org/Vol-410/Paper01.pdf>
9. Bienvenu, M., Bourgaux, C.: Inconsistency-tolerant querying of description logic knowledge bases. In: Pan, J.Z., et al. (eds.) Reasoning Web 2016. LNCS, vol. 9885, pp. 156–202. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-49493-7_5
10. Henzinger, M.R., Henzinger, T.A., Kopke, P.W.: Computing simulations on finite and infinite graphs. In: 36th Annual Symposium on Foundations of Computer Science, Milwaukee, Wisconsin, USA, 23–25 October 1995, pp. 453–462. IEEE Computer Society (1995). <https://doi.org/10.1109/SFCS.1995.492576>
11. Kalyanpur, A., Parsia, B., Horridge, M., Sirin, E.: Finding all justifications of OWL DL entailments. In: Aberer, K., et al. (eds.) ASWC/ISWC -2007. LNCS, vol. 4825, pp. 267–280. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-76298-0_20
12. 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). <https://doi.org/10.1016/j.websem.2015.04.002>

13. Lutz, C., Wolter, F.: Deciding inseparability and conservative extensions in the description logic \mathcal{EL} . *J. Symb. Comput.* **45**(2), 194–228 (2010). <https://doi.org/10.1016/j.jsc.2008.10.007>
14. Matos, V.B., Guimarães, R., Santos, Y.D., Wassermann, R.: Pseudo-contractions as gentle repairs. In: Lutz, C., Sattler, U., Tinelli, C., Turhan, A.-Y., Wolter, F. (eds.) *Description Logic, Theory Combination, and All That*. LNCS, vol. 11560, pp. 385–403. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-22102-7_18
15. Meyer, T.A., Lee, K., Booth, R., Pan, J.Z.: Finding maximally satisfiable terminologies for the description logic \mathcal{ALC} . In: *Proceedings, The Twenty-First National Conference on Artificial Intelligence and the Eighteenth Innovative Applications of Artificial Intelligence Conference*, 16–20 July 2006, Boston, Massachusetts, USA, pp. 269–274. AAAI Press (2006). <http://www.aaai.org/Library/AAAI/2006/aaai06-043.php>
16. Parsia, B., Sirin, E., Kalyanpur, A.: Debugging OWL ontologies. In: *Proceedings of the 14th International Conference on World Wide Web, WWW 2005*, Chiba, Japan, 10–14 May 2005. pp. 633–640. ACM (2005). <https://doi.org/10.1145/1060745.1060837>
17. Peñaloza, R., Turhan, A.-Y.: A practical approach for computing generalization inferences in \mathcal{EL} . In: Antoniou, G., et al. (eds.) *ESWC 2011, Part I*. LNCS, vol. 6643, pp. 410–423. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-21034-1_28
18. Schlobach, S., Cornet, R.: Non-standard reasoning services for the debugging of description logic terminologies. In: *IJCAI-03, Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence, Acapulco, Mexico, 9–15 August 2003*, pp. 355–362. Morgan Kaufmann (2003). <http://ijcai.org/Proceedings/03/Papers/053.pdf>
19. Schlobach, S., Huang, Z., Cornet, R., van Harmelen, F.: Debugging incoherent terminologies. *J. Autom. Reason.* **39**(3), 317–349 (2007). <https://doi.org/10.1007/s10817-007-9076-z>
20. Zarriß, B., Turhan, A.: Most specific generalizations w.r.t. general \mathcal{EL} -TBoxes. In: *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, 3–9 August 2013*, pp. 1191–1197. IJCAI/AAAI (2013). <http://www.aaai.org/ocs/index.php/IJCAI/IJCAI13/paper/view/6709>