

Pushing Optimal ABox Repair from \mathcal{EL} Towards More Expressive Horn-DLs

Franz Baader , Francesco Kriegel

Theoretical Computer Science, Technische Universität Dresden, Dresden, Germany
{franz.baader,francesco.kriegel}@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. DL systems can then be used to compute consequences (such as answers to certain queries) from an ontology consisting of a TBox and an ABox. Since both human-made and machine-learned data sets may contain errors, which manifest themselves as unintuitive or obviously incorrect consequences, repairing DL-based ontologies in the sense of removing such unwanted consequences is an important topic in DL research. Most of the repair approaches described in the literature produce repairs that are not optimal, in the sense that they do not guarantee that only a minimal set of consequences is removed. In a series of papers, we have developed an approach for computing optimal repairs, starting with the restricted setting of an \mathcal{EL} instance store, extending this to the more general setting of a quantified ABox (where some individuals may be anonymous), and then adding a static \mathcal{EL} TBox.

Here, we extend the expressivity of the underlying DL considerably, by adding nominals, inverse roles, regular role inclusions and the bottom concept to \mathcal{EL} , which yields a fragment of the well-known DL Horn- \mathcal{SROIQ} . The ideas underlying our repair approach still apply to this DL, though several non-trivial extensions are needed to deal with the new constructors and axioms. The developed repair approach can also be used to treat unwanted consequences expressed by certain conjunctive queries or regular path queries, and to handle Horn- \mathcal{ALCOI} TBoxes with regular role inclusions.

1 Introduction

Description Logics (DLs) (Baader et al. 2017) 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 The palette of well-investigated DLs with optimized reasoning support goes from the inexpressive and tractable DLs of the \mathcal{EL} and DL-Lite families (Baader, Brandt, and Lutz 2005; Calvanese et al. 2007), on which the OWL 2 profiles OWL 2 EL and OWL 2 QL are based, all the way up to the N2ExpTime-complete DL \mathcal{SROIQ} (Horrocks, Kutz, and Sattler 2006; Kazakov 2008), which is the DL underlying OWL 2. The

consequence-based reasoning approach developed for the \mathcal{EL} family (Baader, Brandt, and Lutz 2005) can be extended to Horn fragments of more expressive DLs, which yields practical “pay as you go” reasoning procedures for these fragments, though they are no longer tractable (Kazakov 2009; Ortiz, Rudolph, and Šimkus 2010).

Like all large human-made digital artefacts, the ontologies employed in applications often contain errors, and this problem is only exacerbated if parts of the ontology (e.g., the data) are automatically generated using inexact methods based on information retrieval or machine learning. Errors are usually detected when reasoning finds an inconsistency or generates consequences that are unintuitive or obviously wrong in the application domain. For the developers of a DL-based ontology it is often quite hard to see how the ontology needs to be modified such that the unwanted consequences no longer follow from the repaired ontology, but as few as possible other consequences are lost.

Classical DL repair approaches based on axiom pinpointing compute maximal subsets of the ontology that do not have the unwanted consequences (Parsia, Sirin, and Kalyanpur 2005; Schlobach et al. 2007; Baader and Suntisrivaraporn 2008). Such repairs depend on the syntactic form of the ontology: if a certain fact is expressed by a single strong axiom rather than an equivalent set of weaker ones, then too many consequences may be lost when removing this strong axiom. To overcome this problem, more fine-grained approaches for repairing DL-based ontologies have been developed (Horridge, Parsia, and Sattler 2008; Lam et al. 2008; Du, Qi, and Fu 2014; Troquard et al. 2018; Baader et al. 2018). These approaches are, however, still not optimal since they apply some restrictions on how the ontology can be changed, based on its syntactic form. In particular, they usually do not add new objects to the ABox.

To see why new objects may be needed to achieve optimality, assume that the ABox contains the information that Kim, who is rich and famous, is Ann’s child, expressed by the assertions $Famous(KIM)$, $Rich(KIM)$, and $child(ANN, KIM)$, and that we want to remove the consequence $\exists child.(Rich \sqcap Famous)(ANN)$. If we decide to keep the assertion that Kim is Ann’s child, then we need to remove either $Rich(KIM)$ or $Famous(KIM)$. However, if we decide that this Kim is not Ann’s child after all, simply removing the role assertion $child(ANN, KIM)$ would

¹<https://www.w3.org/TR/owl2-overview/>

also remove implied consequences for Ann. This can be avoided by adding the assertions $child(ANN, x)$, $Rich(x)$, $child(ANN, y)$, and $Famous(y)$, where x and y are anonymous individuals, which are formally represented in a quantified ABox (qABox) by existentially quantified variables. This example illustrates the main idea underlying the optimal repair approach introduced in (Baader et al. 2020): the use of quantified ABoxes and the construction of appropriate anonymous copies of individuals. The main technical problem to solve in (Baader et al. 2020) was to find out which copies with what properties are needed to achieve optimality. This work dealt with an input ontology consisting only of a qABox, and assumed that the unwanted consequences are instance relationships $C(a)$ for \mathcal{EL} concepts C .

In (Baader et al. 2021), we extended this approach to a setting where, in addition to the qABox, the ontology contains an \mathcal{EL} TBox, which is assumed to be correct, and thus cannot be changed during repair. To add consequences implied by the TBox, we saturate the qABox by using the concept inclusions as rewrite rules before repairing the qABox. If, in our example, the TBox contained the concept inclusion $Celebrity \sqsubseteq Rich \sqcap Famous$ and the ABox contained $Celebrity(KIM)$ rather than $Rich(KIM)$ and $Famous(KIM)$, then saturation would add the latter two assertions. If then $Celebrity(KIM)$ is removed in the repair, these two consequences can still be preserved. However, when repairing the saturated qABox, care must be taken that the TBox cannot re-introduce assertions that have been removed by the repair. For example, in the case where the unwanted consequence is $Rich(KIM)$, it is not enough to remove this assertion from the saturated qABox: one also needs to remove $Celebrity(KIM)$ since together with the TBox it implies $Rich(KIM)$. The problem with saturation is that, in the presence of cyclic concept inclusions, such as $Rich \sqsubseteq \exists child.Rich$, it may not terminate. This is not just a problem of our repair approach, but may prevent the existence of optimal repairs (see Example 9 below). In (Baader et al. 2021), two approaches are considered to overcome this problem. On the one hand, one can restrict the attention to TBoxes that are cycle-restricted as introduced in (Baader, Borgwardt, and Morawska 2012). On the other hand, if one is only interested in answers to instance queries, one can apply a weaker saturation operation, called IQ-saturation, which always terminates for \mathcal{EL} .

In this paper, we extend the expressivity of the DL used to formulate the TBox and the unwanted consequences considerably, by adding nominals, inverse roles, role inclusion axioms, and the bottom concept to \mathcal{EL} . To obtain a decidable DL and guarantee the existence of optimal repairs, we restrict the set of role inclusion axioms to being regular, as in the DL $SROIQ$. In addition, we first consider the case without the bottom concept, and only later deal with the additional problems caused by the fact that bottom may cause the ontology to become inconsistent. Computability of the set of optimal repairs and the fact that this set covers all repairs (in the sense that every repair is entailed by an optimal one) follows from a “small repair” property, which can be shown using an adaptation of the well-known filtration technique (Baader et al. 2017). However, even disregarding

the impracticality of an algorithm that computes the optimal repairs by looking at all qABoxes up to a certain size bound, this does not lead to a viable method for choosing an appropriate optimal repair since it would require the knowledge engineer to choose among exponentially many repairs of exponential size. In contrast, the *canonical repairs* (which cover all optimal repairs) constructed by our extension of the repair approach in (Baader et al. 2021) are characterized by so-called *repair seeds*, which are of polynomial size. The knowledge engineer can choose among these by answering a polynomial number of instance queries (i.e., queries about which instance relationships hold in the application domain).

The added expressivity generates new challenges, which require non-trivial adaptations of our approach for constructing canonical repairs. Since nominals in the TBox can imply equality between individuals, we extend qABoxes by equality assertions, to be able to represent such consequences in the saturated qABox, and we also must repair unwanted equalities. We deal with role inclusion axioms and inverse roles by using finite automata, which can represent the infinitely many implied role inclusions in a finite way. Technically, this is where we make use of the restriction to regular sets of role inclusion axioms. To handle inconsistency caused by bottom, we consider not only “local” unwanted consequences of the form $C(a)$, but also “global” ones of the form $\exists\{x\}.C(x)$. If, in our example, the TBox additionally says that rich and poor are disjoint, using the concept inclusion $Poor \sqcap Rich \sqsubseteq \perp$, and the qABox states that Ann has an (anonymous) child that is a poor celebrity, then the entailed inconsistency can be repaired by preventing the consequence $\exists\{x\}.((Poor \sqcap Rich)(x))$.

The added expressivity also allows us to specify interesting kinds of unwanted consequences other than instance relationships. On the one hand, we can deal with regular reachability queries, which are similar to regular path queries (Calvanese, Eiter, and Ortiz 2009). On the other hand, we can also treat certain kinds of conjunctive queries. The problem of repairing w.r.t. conjunctive queries to qABoxes has already been considered, in the guise of achieving compliance for relational datasets with labelled nulls, in (Grau and Kostylev 2019). However, this work does not allow for background TBoxes, and the notion of optimality used there is different from ours since it restricts the possible changes to the qABox to a sequence of certain anonymization operations. Finally, our repair approach can also deal with Horn- $ALCOI$ -TBoxes together with sets of regular role inclusion axioms.

Due to space restrictions, full proofs of our results are given in (Baader and Kriegel 2022).

2 Preliminaries

First, we introduce the DL \mathcal{ELROI} employed to formulate terminological background knowledge, and the quantified ABoxes with equalities used to represent the data. Then we describe how such an ABox can be saturated w.r.t. the terminological knowledge, and finally define regular sets of role inclusions and show how to represent them using finite automata.

2.1 The Description Logic \mathcal{ELROI}

The DL \mathcal{ELROI} extends \mathcal{EL} with (complex) role inclusions (\mathcal{R}), nominals (\mathcal{O}), and inverse roles (\mathcal{I}). Let Σ be a *signature*, i.e., a disjoint union of finite, non-empty sets Σ_I , Σ_C , and Σ_R of *individual names*, *concept names*, and *role names*, respectively. A *role* is either a role name or an *inverse role* r^- for some role name $r \in \Sigma_R$. For a role R we write R^- to denote r^- if $R = r$ is a role name and r if $R = r^-$ is an inverse role. Concept descriptions C of \mathcal{ELROI} are constructed using the grammar rule

$$C ::= \top \mid A \mid \{a\} \mid C \sqcap C \mid \exists R.C,$$

where A ranges over concept names, a over individual names, and R over roles. An *atom* is a concept name A , a *nominal* $\{a\}$, or an *existential restriction* $\exists R.C$. Each concept description C is a conjunction of atoms, with \top corresponding to the empty conjunction. We denote the set of these atoms as $\text{Conj}(C)$.

A *concept inclusion* (CI) is of the form $C \sqsubseteq D$ for concept descriptions C, D , and a *role inclusion* (RI) is of the form $\varepsilon \sqsubseteq S$ or $R_1 \circ \dots \circ R_n \sqsubseteq S$ for roles R_1, \dots, R_n, S and $n \geq 1$. In the following, when we write $R_1 \circ \dots \circ R_n \sqsubseteq S$, we assume that $n \geq 0$, where $R_1 \circ \dots \circ R_n$ for $n = 0$ stands for ε . A *TBox* is a finite set of CIs, an *RBox* is a finite set of RIs, and a pair $(\mathcal{T}, \mathcal{R})$ consisting of a TBox \mathcal{T} and an RBox \mathcal{R} is called a *terminology*. A *concept assertion* $C(a)$ is a shorthand for the CI $\{a\} \sqsubseteq C$, and a *role assertion* $r(a, b)$ abbreviates $\{a\} \sqsubseteq \exists r.\{b\}$. Furthermore, $r^-(a, b)$ means $r(b, a)$.

The semantics of \mathcal{ELROI} is defined as usual (Baader et al. 2017; Horrocks, Kutz, and Sattler 2006) using the notion of an *interpretation* $\mathcal{I} = (\text{Dom}(\mathcal{I}), \cdot^{\mathcal{I}})$, which assigns subsets $C^{\mathcal{I}}$ of $\text{Dom}(\mathcal{I})$ to concepts C and binary relations $R^{\mathcal{I}}$ on $\text{Dom}(\mathcal{I})$ to roles R according to the semantics of the constructors. Individual names a are mapped to elements $a^{\mathcal{I}}$ of $\text{Dom}(\mathcal{I})$, without requiring the *unique name assumption*, i.e., $a^{\mathcal{I}} = b^{\mathcal{I}}$ is allowed for distinct individual names a, b . Models of TBoxes and RBoxes are also defined in the usual way. We say that the terminology $(\mathcal{T}, \mathcal{R})$ *entails* a CI or RI α (written $(\mathcal{T}, \mathcal{R}) \models \alpha$) if α holds in every model of \mathcal{T} and \mathcal{R} . In case $(\mathcal{T}, \mathcal{R}) \models C \sqsubseteq D$ we say that C is *subsumed* by D w.r.t. $(\mathcal{T}, \mathcal{R})$, and may write $C \sqsubseteq^{\mathcal{T}, \mathcal{R}} D$ to express this.

Note that other interesting axioms concerning roles can be expressed using RIs and inverse roles. Reflexivity, transitivity, and symmetry of r can respectively be enforced by the RIs $\varepsilon \sqsubseteq r$, $r \circ r \sqsubseteq r$, and $r \sqsubseteq r^-$, and *range restrictions* $\text{Ran}(r) \sqsubseteq C$ can be expressed by CIs $\exists r^-. \top \sqsubseteq C$.

This last observation shows that subsumption in \mathcal{ELROI} is actually undecidable since it was shown in (Baader, Lutz, and Brandt 2008) that subsumption in \mathcal{EL} w.r.t. RIs and range restrictions is undecidable. We will avoid this problem by imposing a restriction on RBoxes (see Section 2.4).

2.2 Quantified ABoxes with Equalities

Quantified ABoxes were first introduced in (Baader et al. 2020), but they were also considered, as relational datasets with labelled nulls, in (Grau and Kostylev 2019), and their

existentially quantified variables correspond to the “anonymous individuals” in the OWL2 standard. Also, as explained in (Baader et al. 2020), quantified ABoxes are basically the same as Boolean conjunctive queries. Here, we extend this notion by allowing for equality assertions, but for simplicity still use the name “quantified ABoxes” for the extended formalism. Equality assertions are used to represent implied equality between individuals; e.g., the CI $\{a\} \sqsubseteq \{b\}$ implies that a and b must always be interpreted by the same element of the domain.

Let Σ be a signature. A *quantified ABox* ($qABox$) $\exists X.\mathcal{A}$ over Σ consists of a finite set X of *variables*, which is disjoint with Σ , and a *matrix* \mathcal{A} , which is a finite set of *concept assertions* $A(u)$, *role assertions* $r(u, v)$, and *equality assertions* $a \equiv b$, where $A \in \Sigma_C$, $r \in \Sigma_R$, $u, v \in \Sigma_I \cup X$, and $a, b \in \Sigma_I$. An *object name* of $\exists X.\mathcal{A}$ is either an element of Σ_I or a variable in X . We denote the set of these objects as $\text{Obj}(\exists X.\mathcal{A})$. If X is empty, then we sometimes drop the quantifier $\exists \emptyset$.

The interpretation \mathcal{I} is a *model* of $\exists X.\mathcal{A}$ (written $\mathcal{I} \models \exists X.\mathcal{A}$) if there is a *variable assignment* $\mathcal{Z}: X \rightarrow \text{Dom}(\mathcal{I})$ such that the augmented interpretation $\mathcal{I}[\mathcal{Z}]$ that additionally maps each variable x to $\mathcal{Z}(x)$ is a model of the matrix \mathcal{A} , i.e., $u^{\mathcal{I}[\mathcal{Z}]} \in A^{\mathcal{I}}$ for each $A(u) \in \mathcal{A}$, $(u^{\mathcal{I}[\mathcal{Z}]}, v^{\mathcal{I}[\mathcal{Z}]}) \in r^{\mathcal{I}}$ for each $r(u, v) \in \mathcal{A}$, and $a^{\mathcal{I}} = b^{\mathcal{I}}$ for each $a \equiv b \in \mathcal{A}$. Given a terminology $(\mathcal{T}, \mathcal{R})$ and $qABoxes$ $\exists X.\mathcal{A}$ and $\exists Y.\mathcal{B}$, we say that $\exists X.\mathcal{A}$ *entails* $\exists Y.\mathcal{B}$ w.r.t. $(\mathcal{T}, \mathcal{R})$ (written $\exists X.\mathcal{A} \models^{\mathcal{T}, \mathcal{R}} \exists Y.\mathcal{B}$) if every model of $\exists X.\mathcal{A}$ and $(\mathcal{T}, \mathcal{R})$ is also a model of $\exists Y.\mathcal{B}$. If both the TBox \mathcal{T} and the RBox \mathcal{R} are empty, then we omit the suffix “w.r.t. $(\mathcal{T}, \mathcal{R})$ ” and write \models instead of $\models^{\mathcal{T}, \mathcal{R}}$. Similar simplifications are made if one of them is empty.

For $qABoxes$ without equality assertions, it was shown in (Baader et al. 2020) that entailment can be characterized using homomorphisms. In our extended setting, we need to adapt the definition of a homomorphism between $qABoxes$. To this purpose, we consider the equivalence relation $\approx_{\exists X.\mathcal{A}}$ on $\text{Obj}(\exists X.\mathcal{A})$ induced by the equality assertions in $\exists X.\mathcal{A}$, which is defined as the reflexive, symmetric, transitive closure of the relation $\{(a, b) \mid a \equiv b \in \mathcal{A}\}$. We sometimes write \approx for $\approx_{\exists X.\mathcal{A}}$ if the $qABox$ is clear from the context, and denote the equivalence classes by $[u]_{\exists X.\mathcal{A}}$. Since there are no equality assertions involving variables, each equivalence class of a variable is a singleton set.

Definition 1. A *homomorphism* h from a $qABox$ $\exists X.\mathcal{A}$ to a $qABox$ $\exists Y.\mathcal{B}$ is a mapping $h: \text{Obj}(\exists X.\mathcal{A}) \rightarrow \text{Obj}(\exists Y.\mathcal{B})$ that satisfies the following conditions:

- (Hom1)** $a \approx_{\exists X.\mathcal{A}} b$ implies $h(a) \approx_{\exists Y.\mathcal{B}} h(b)$ for all individual names a, b .
- (Hom2)** $h(a) = a$ for each individual name a .
- (Hom3)** For each concept assertion $A(t) \in \mathcal{A}$, there is an object name v such that $v \approx_{\exists Y.\mathcal{B}} h(t)$ and $A(v) \in \mathcal{B}$.
- (Hom4)** For each role assertion $r(t, u) \in \mathcal{A}$, there are object names v, w such that $v \approx_{\exists Y.\mathcal{B}} h(t)$, $w \approx_{\exists Y.\mathcal{B}} h(u)$, and $r(v, w) \in \mathcal{B}$.

Based on this notion of homomorphism, entailment between $qABoxes$ with equality assertions can now be characterized as follows.

Conjunction Rule. If \mathcal{B} contains the assertion $(C_1 \sqcap \dots \sqcap C_n)(t)$ for $n \neq 1$, then remove it from \mathcal{B} and add the assertions $C_1(t), \dots, C_n(t)$ to \mathcal{B} .

Existential Restriction Rule. If \mathcal{B} contains the assertion $\exists R.C(t)$, then remove it from \mathcal{B} , add a fresh variable y to Y , and add the assertions $R(t, y)$ and $C(y)$ to \mathcal{B} .

Nominal Rule. If \mathcal{B} contains the assertion $\{a\}(t)$, then remove it from \mathcal{B} and, if t is an individual name, then add the equality $t \equiv a$ to \mathcal{B} ; otherwise replace every occurrence of t in \mathcal{B} by a and remove t from Y .

Concept Inclusion Rule. If \mathcal{T} contains the CI $C \sqsubseteq D$ and \mathcal{B} entails the concept assertion $C(t)$, but not $D(t)$, then add the concept assertion $D(t)$ to \mathcal{B} .

Role Inclusion Rule. If \mathcal{R} contains the RI $R_1 \circ \dots \circ R_n \sqsubseteq S$ and \mathcal{B} entails the role assertions $R_1(t_0, t_1), \dots, R_n(t_{n-1}, t_n)$, but not $S(t_0, t_n)$, then add the role assertion $S(t_0, t_n)$ to \mathcal{B} .

Figure 1: The saturation rules are exhaustively applied to a qABox $\exists Y.\mathcal{B}$ w.r.t. a terminology $(\mathcal{T}, \mathcal{R})$, starting with $\exists Y.\mathcal{B} := \exists X.\mathcal{A}$ for an input qABox $\exists X.\mathcal{A}$.

Proposition 2. *The qABox $\exists X.\mathcal{A}$ is entailed by the qABox $\exists Y.\mathcal{B}$ iff there exists a homomorphism from $\exists X.\mathcal{A}$ to $\exists Y.\mathcal{B}$.*

As in the case of qABoxes without equality assertions, this provides us with an NP decision procedure for entailment. NP-hardness already holds without equality assertions (Baader et al. 2020).

We often need to consider the matrix \mathcal{A} of a quantified ABox $\exists X.\mathcal{A}$ alone, without the quantifier prefix. We can view \mathcal{A} to be an “ordinary” ABox without quantifiers (or equivalently as a qABox with empty quantifier prefix) by extending the signature to $\Sigma \cup X$, where variables are treated as individuals. This allows us to evaluate entailment expressions like $\mathcal{A} \models C(x)$, where C is a concept description and $x \in X$, using interpretations and models for the extended signature.

2.3 Saturation

The purpose of saturation is to extend a given qABox $\exists X.\mathcal{A}$ with enough consequences derived using the terminology $(\mathcal{T}, \mathcal{R})$ such that entailment from $\exists X.\mathcal{A}$ w.r.t. $(\mathcal{T}, \mathcal{R})$ is the same as entailment from its saturation $\text{sat}^{\mathcal{T}, \mathcal{R}}(\exists X.\mathcal{A}) = \exists Y.\mathcal{B}$ w.r.t. the empty terminology. The rules in Figure 1 extend the CQ-saturation rules in (Baader et al. 2021) such that nominals, inverse roles, and RIs are taken into account. Note that, during saturation, the matrix \mathcal{B} may contain complex concept assertions, but after termination all concept assertions are again restricted to concept names. The semantics of qABoxes with complex concept assertions is defined in the obvious way.

In general, application of the saturation rules need not terminate, already in the \mathcal{EL} setting considered in (Baader et al. 2021). But there the restriction to cycle-restricted TBoxes guarantees termination, where an \mathcal{EL} TBox \mathcal{T} is cycle-restricted if there is no concept C and roles r_1, \dots, r_n

($n \geq 1$) such that $C \sqsubseteq_{\mathcal{T}} \exists r_1. \dots \exists r_n.C$. For \mathcal{ELROI} terminologies, the RBox may cause non-termination even if the TBox is cycle-restricted.

Example 3. Consider the \mathcal{ELROI} TBox $\mathcal{T} := \{A \sqsubseteq \exists r.\top, \exists s.\top \sqsubseteq \exists s.A\}$, the RBox $\mathcal{R} := \{r \sqsubseteq s\}$, and the qABox $\exists \emptyset.\mathcal{A}$ with $\mathcal{A} := \{A(a)\}$. The TBox \mathcal{T} is cycle-restricted and saturation of $\exists \emptyset.\mathcal{A}$ with (\mathcal{T}, \emptyset) terminates after a has received an r -successor x_1 . However, w.r.t. $(\mathcal{T}, \mathcal{R})$, the role inclusion rule makes x_1 also an s -successor of a . The concept inclusion rule then adds an s -successor y_1 of a and the assertion $A(y_1)$. But now y_1 receives an r -successor x_2 , which becomes an s -successor of y_1 , etc.

Since our repair approach works on saturated qABoxes, it can only be applied in the presence of terminologies $(\mathcal{T}, \mathcal{R})$ that are terminating in the following sense.

Definition 4. The terminology $(\mathcal{T}, \mathcal{R})$ is *terminating* if, for each qABox $\exists X.\mathcal{A}$, there is a finite sequence of applications of the saturation rules in Figure 1 to $\exists X.\mathcal{A}$ resulting in a qABox to which no more rule applies. We then denote this qABox as $\text{sat}^{\mathcal{T}, \mathcal{R}}(\exists X.\mathcal{A})$ and call it the *saturation* of $\exists X.\mathcal{A}$ w.r.t. $(\mathcal{T}, \mathcal{R})$.

We refrain here from giving our own decidable sufficient condition for termination of a terminology $(\mathcal{T}, \mathcal{R})$. Instead, we point out that one can translate the concept inclusions in \mathcal{T} and the role inclusions in \mathcal{R} into a set of existential rules, and that saturation then corresponds to applying the so-called chase. One can thus try to use one of the numerous acyclicity conditions guaranteeing chase termination proposed in the database and rules communities (see, e.g., (Grau et al. 2013)) to show termination of $(\mathcal{T}, \mathcal{R})$. The saturation obtained in case of termination has the following important property.

Theorem 5. *Let $(\mathcal{T}, \mathcal{R})$ be a terminating terminology and $\exists X.\mathcal{A}$ a quantified ABox. Then, for every qABox $\exists Z.\mathcal{C}$, the following statements are equivalent:*

1. $\exists X.\mathcal{A} \models^{\mathcal{T}, \mathcal{R}} \exists Z.\mathcal{C}$.
2. $\text{sat}^{\mathcal{T}, \mathcal{R}}(\exists X.\mathcal{A}) \models \exists Z.\mathcal{C}$.
3. *There is a homomorphism from $\exists Z.\mathcal{C}$ to $\text{sat}^{\mathcal{T}, \mathcal{R}}(\exists X.\mathcal{A})$.*

In (Baader et al. 2021), a different kind of saturation, called IQ-saturation, was introduced, which always terminates. Using IQ-saturation in the repair process was shown to be sufficient if one is only interested in instance queries. However, due to the presence of inverse roles in \mathcal{ELROI} , it is easy to see that finite IQ-saturations cannot always work (see (Baader and Kriegel 2022) for an example).

2.4 Regular RBoxes

As pointed out at the end of Section 2.1, subsumption is undecidable in \mathcal{ELROI} if arbitrary RBoxes are allowed. In (Baader, Lutz, and Brandt 2008), tractability of \mathcal{EL}^{++} is ensured by restricting the interaction between range restrictions and RIs. Since, in our setting, range restrictions are expressed using inverse roles and CIs, it is not clear how to adapt this solution. Instead, we use the regularity restriction imposed in (Horrocks, Kutz, and Sattler 2006; Kazakov 2008) to make \mathcal{SROIQ} decidable, which is required by our repair approach anyway.

Definition 6. An RBox \mathcal{R} is *regular* if, for each role R , the language $L_{\mathcal{R}}(R) := \{S_1 \cdots S_n \mid S_1 \circ \cdots \circ S_n \sqsubseteq^{\mathcal{R}} R\}$ is regular. The sublogic of \mathcal{ELROI} that only supports regular RBoxes is denoted by $\mathcal{ELR}_{\text{reg}}\mathcal{OI}$.

Since $\mathcal{ELR}_{\text{reg}}\mathcal{OI}$ is a fragment of Horn- \mathcal{SROIQ} , it inherits the complexity upper-bound of 2ExpTime (Ortiz, Rudolph, and Šimkus 2010). The exact complexity of subsumption in $\mathcal{ELR}_{\text{reg}}\mathcal{OI}$ is open, with the best lower-bound of ExpTime inherited from \mathcal{ELI} (Baader, Lutz, and Brandt 2008).

To the best of our knowledge, it is not known whether RBox regularity is decidable. Decidability of the closely related regularity problem for pure context-free grammars has been open for a long time (Maurer, Salomaa, and Wood 1980). However, there exist syntactic restrictions that guarantee regularity (Horrocks, Kutz, and Sattler 2006; Kazakov 2010), and if these restrictions apply then one can effectively construct (exponentially large) finite automata accepting the regular languages $L_{\mathcal{R}}(R)$.

Let \mathcal{R} be a regular RBox, and for each role R , let $\mathfrak{A}_R = (Q_R, \Sigma_R^{\pm}, i_R, \Delta_R, F_R)$ be a finite automaton (with set of states Q_R , the alphabet Σ_R^{\pm} of all roles, initial state i_R , transition relation Δ_R , and set of final states F_R) accepting $L_{\mathcal{R}}(R)$, i.e., such that $L(\mathfrak{A}_R) = L_{\mathcal{R}}(R)$. We assume without loss of generality (but in the worst-case paid for by another exponential blowup) that each automaton \mathfrak{A}_R is deterministic. In addition, we assume that \mathfrak{A}_R does not contain states that are unreachable from the initial state or from which no final state can be reached, and further that the sets Q_R for different R are pairwise disjoint and are all disjoint with the signature Σ . For each state $q \in Q_R$, the automaton $\mathfrak{A}_{\mathcal{R}}(q) := (Q_R, \Sigma_R^{\pm}, q, \Delta_R, F_R)$ is obtained from \mathfrak{A}_R by replacing the initial state i_R with q . We will use existential restrictions of the form $\exists q.C$ for $q \in Q_R$ as abbreviations for the (possibly infinite) disjunction $\bigsqcup \{ \exists S_1 \cdots \exists S_n.C \mid S_1 \cdots S_n \in L(\mathfrak{A}_{\mathcal{R}}(q)) \}$, i.e., in each interpretation \mathcal{I} , $(\exists q.C)^{\mathcal{I}}$ is defined to be

$$\bigcup \{ (\exists S_1 \cdots \exists S_n.C)^{\mathcal{I}} \mid S_1 \cdots S_n \in L(\mathfrak{A}_{\mathcal{R}}(q)) \}.$$

Entailment for such existential restrictions can be characterized as follows.

Lemma 7. *Given a qABox $\exists X.A$, an object t of it, a terminology $(\mathcal{T}, \mathcal{R})$ with regular \mathcal{R} , \mathcal{ELROI} concept descriptions C, D , and a state q . Then the following holds:*

1. $\mathcal{A} \models^{\mathcal{T}, \mathcal{R}} \exists q.C(t)$ iff there is a word $S_1 \cdots S_n \in L(\mathfrak{A}_{\mathcal{R}}(q))$ such that $\mathcal{A} \models^{\mathcal{T}, \mathcal{R}} \exists S_1 \cdots \exists S_n.C(t)$,
2. $D \sqsubseteq^{\mathcal{T}, \mathcal{R}} \exists q.C$ iff there is a word $S_1 \cdots S_n \in L(\mathfrak{A}_{\mathcal{R}}(q))$ such that $D \sqsubseteq^{\mathcal{T}, \mathcal{R}} \exists S_1 \cdots \exists S_n.C$.

If the terminology is terminating, we can decide whether the conditions for entailment stated in Lemma 7 hold. Basically, to check whether $\mathcal{A} \models^{\mathcal{T}, \mathcal{R}} \exists q.C(t)$ holds, we simply need to find an accepting run of the automaton $\mathfrak{A}_{\mathcal{R}}(q)$ such that the accepted word corresponds to a path in the saturation $\text{sat}^{\mathcal{T}, \mathcal{R}}(\exists X.A)$ that starts with t and ends with an instance of C . This boils down to a reachability test in the product of the automaton with the saturation. The second condition in Lemma 7 can be reduced to the first one.

3 Optimal and Canonical Repairs

In this section, we first extend the notion of an (optimal) repair, as introduced in (Baader et al. 2021), to the more expressive DL \mathcal{ELROI} and a setting where the repair request, which describes which consequences are to be removed, also contains global unwanted consequences. For regular sets of role inclusions, we show that every repair is entailed by a repair containing a bounded number of individuals. From this, we derive that the set of optimal repairs can effectively be computed and covers all repairs. Then, we extend the construction of canonical repairs of (Baader et al. 2021) from \mathcal{EL} to $\mathcal{ELR}_{\text{reg}}\mathcal{OI}$. The set of canonical repairs can effectively be computed, covers all repairs and thus contains all optimal repairs, and a repair seed determining such a canonical repair can be chosen by answering a polynomial number of instance queries. Throughout the section, we assume (unless specified otherwise) that $\exists X.A$ is a quantified ABox, \mathcal{T} an \mathcal{ELROI} TBox, \mathcal{R} a regular RBox, all defined over the same signature Σ , and that $(\mathcal{T}, \mathcal{R})$ is terminating.

Definition 8. A *repair request* \mathcal{P} is a union of a finite set \mathcal{P}_{loc} of \mathcal{ELROI} concept assertions, the *local* request, and of a finite set \mathcal{P}_{glo} of \mathcal{ELROI} concept descriptions, the *global* request. A *repair* of $\exists X.A$ for \mathcal{P} w.r.t. $(\mathcal{T}, \mathcal{R})$ is a quantified ABox $\exists Y.B$ that fulfills the following properties:

(Rep1) $\exists X.A \models^{\mathcal{T}, \mathcal{R}} \exists Y.B$,

(Rep2) $\exists Y.B \not\models^{\mathcal{T}, \mathcal{R}} C(a)$ for each $C(a) \in \mathcal{P}_{\text{loc}}$,

(Rep3) $\exists Y.B \not\models^{\mathcal{T}, \mathcal{R}} \exists \{x\}. \{D(x)\}$ for each $D \in \mathcal{P}_{\text{glo}}$.

This repair is *optimal* if there is no repair $\exists Z.C$ such that $\exists Z.C \models^{\mathcal{T}, \mathcal{R}} \exists Y.B$, but $\exists Y.B \not\models^{\mathcal{T}, \mathcal{R}} \exists Z.C$. We say that a set of repairs \mathfrak{S} *covers* all repairs if every repair is entailed w.r.t. $(\mathcal{T}, \mathcal{R})$ by a repair in \mathfrak{S} .

Obviously, $\exists X.A$ has a repair for \mathcal{P} w.r.t. $(\mathcal{T}, \mathcal{R})$ iff the terminology alone does not imply any of the unwanted consequences in \mathcal{P} , since then the empty ABox is a repair. The restriction to terminating terminologies and regular RBoxes is needed to ensure that any repair problem has a finite set of optimal repairs covering all repairs. The proof of Proposition 2 in (Baader et al. 2018) contains an example with non-terminating terminology where there is no optimal repair, though there is a repair. However, in this proof it is only shown that there cannot be an optimal repair that is an ABox. While this proof can be adapted to deal also with qABoxes, we present here a modified example with exactly one optimal repair, which however does not cover all repairs.

Example 9. Assume that $\mathcal{T} := \{A \sqsubseteq \exists r.A, \exists r.A \sqsubseteq A\}$, $\mathcal{R} := \emptyset$, $\mathcal{A} := \{A(a), B(a)\}$, and $\mathcal{P} := \{(A \sqcap B)(a)\}$. Then $\exists \{x\}. \{A(a), A(x), B(x)\}$ is an optimal repair of $\exists \emptyset.A$ for \mathcal{P} w.r.t. $(\mathcal{T}, \mathcal{R})$. However, there are also repairs in which the concept assertion $B(a)$ is retained, and $A(a)$ is removed. To see that there cannot be an optimal repair containing $B(a)$, first note that \mathcal{A} together with \mathcal{T} does not imply the existence of any role cycle, and thus no repair can contain such a cycle. Consequently, for an optimal repair $\exists Y.B$ containing $B(a)$, there is an upper bound n on the length of role chains starting from a . Adding $r(a, y_1), r(y_2, y_3), \dots, r(y_n, y_{n+1})$ for fresh existentially quantified variables y_1, \dots, y_{n+1} to

$\exists Y.B$ then yields a new repair that strictly implies $\exists Y.B$, which contradicts the assumed optimality of this repair.

The following example shows that non-regularity of the RBox may prevent all repairs from being covered by a finite set of repairs.

Example 10. The RBox $\mathcal{R} := \{r^- \circ s \circ r \sqsubseteq s\}$ is not regular since $L_{\mathcal{R}}(s) = \{(r^-)^i s r^i \mid i \geq 0\}$ is a context-free language over the alphabet $\{r^-, s, r\}$ known to be non-regular. Together with the TBox $\mathcal{T} := \{\exists s.A \sqsubseteq A, \exists s.B \sqsubseteq B\}$, this RBox yields a terminating terminology. Consider the ABox $\mathcal{A} := \{r(a, a), s(a, a), A(a), B(a)\}$ and the repair request $\mathcal{P} = \mathcal{P}_{\text{glo}} := \{A \sqcap B\}$. It is not hard to see that, for each $n \geq 1$, the qABox $\exists X_n.A_n$ is a repair of $\exists \emptyset.A$ for \mathcal{P} w.r.t. $(\mathcal{T}, \mathcal{R})$, where $X_n := \{x_1, \dots, x_n\}$ and

$$\begin{aligned} \mathcal{A}_n := \{ & r(a, x_1), r(x_1, x_2), \dots, r(x_{n-1}, x_n), \\ & s(a, a), s(x_1, x_1), \dots, s(x_n, x_n), \\ & A(a), A(x_1), A(x_2), \dots, A(x_{n-1}), B(x_n)\}. \end{aligned}$$

Assume that \mathfrak{S} is a finite set of repairs of $\exists \emptyset.A$ for \mathcal{P} w.r.t. $(\mathcal{T}, \mathcal{R})$ that covers all repairs, and let n be larger than the maximal number of objects occurring in the elements of \mathfrak{S} . Without loss of generality we assume that the elements of \mathfrak{S} are saturated w.r.t. $(\mathcal{T}, \mathcal{R})$. Then there must exist a repair $\exists Y.B$ in \mathfrak{S} such that there is a homomorphism h from $\exists X_n.A_n$ to $\exists Y.B$. Since B contains less than n objects, there must be i, j with $1 \leq i < j \leq n$ such that $h(x_i) = h(x_j)$. Consequently, $h(x_n)$ is reachable from $h(a)$ with the role r both in n steps and in $m < n$ steps, where $m = n - (j - i)$. Since $h(x_m)$ is also reachable in m steps from $h(a)$ and $s(h(a), h(a))$ must be in B , the fact that $\exists Y.B$ is saturated implies that $s(h(x_n), h(x_m))$ must belong to B . Since $A(x_m) \in \mathcal{A}_n$ yields $A(h(x_m)) \in B$, this implies that $A(h(x_n)) \in B$. However, since $B(x_n) \in \mathcal{A}_n$ also yields $B(h(x_n)) \in B$, this contradicts our assumption that $\exists Y.B$ is a repair for \mathcal{P} .

The Small Repair Property If we restrict the attention to terminating terminologies with regular RBoxes \mathcal{R} , then we can show that the repairs of a certain bounded size cover all repairs. For an \mathcal{ELROI} TBox \mathcal{T} and a repair request \mathcal{P} , let $\text{Sub}(\mathcal{T}, \mathcal{P})$ denote the set of concept descriptions occurring in \mathcal{T} and \mathcal{P} and $\text{Atoms}(\mathcal{T}, \mathcal{P})$ the set of atoms in this set. To take the RBox into account, we introduce the set of \mathcal{R} -extended atoms $\text{Atoms}_{\mathcal{R}}(\mathcal{T}, \mathcal{P})$, which is obtained from $\text{Atoms}(\mathcal{T}, \mathcal{P})$ by replacing each $\exists R.C \in \text{Atoms}(\mathcal{T}, \mathcal{P})$ with the existential restrictions $\exists q.C$, where q ranges over $Q_{\mathcal{R}}$ (i.e., the set of states of the automaton for $L_{\mathcal{R}}(R)$).

Proposition 11. *Let $(\mathcal{T}, \mathcal{R})$ be a terminating \mathcal{ELROI} terminology with regular RBox, \mathcal{P} an \mathcal{ELROI} repair request, $\exists X.A$ a (w.l.o.g) saturated qABox with m objects, and $n := |\text{Atoms}(\mathcal{T}, \mathcal{P}) \cup \text{Atoms}_{\mathcal{R}}(\mathcal{T}, \mathcal{P})|$. Then every repair of $\exists X.A$ for \mathcal{P} w.r.t. $(\mathcal{T}, \mathcal{R})$ is entailed w.r.t. $(\mathcal{T}, \mathcal{R})$ by a repair that contains at most $m \cdot 2^n$ objects.*

This proposition can be shown by adapting the well-known filtration technique, e.g., used in (Baader et al. 2017) to prove the finite model property for \mathcal{ALC} . Let $\exists Y.B$ be a repair of $\exists X.A$ for \mathcal{P} w.r.t. $(\mathcal{T}, \mathcal{R})$, and assume without loss of generality that it is saturated. Since $\exists X.A$ entails every repair and is also assumed to be saturated, there

is a homomorphism h from $\exists Y.B$ to $\exists X.A$. For objects u of $\exists Y.B$, we set $t(u) := \{C \mid C \in \text{Atoms}(\mathcal{T}, \mathcal{P}) \cup \text{Atoms}_{\mathcal{R}}(\mathcal{T}, \mathcal{P}) \text{ and } B \models C(u)\}$, and define the equivalence relation \sim on these objects as

$$u \sim v \text{ iff } t(u) = t(v) \text{ and } h(u) = h(v).$$

Obviously, \sim has at most $m \cdot 2^n$ equivalence classes $[u]_{\sim}$. The filtration $\exists Z.C$ has these equivalence classes as objects, with the class $[a]_{\sim}$ standing for the individual a . The classes inherit their concept and role assertions from the ones of their elements in B (see (Baader and Kriegel 2022) for details). We can then show, for all $C \in \text{Atoms}(\mathcal{T}, \mathcal{P}) \cup \text{Atoms}_{\mathcal{R}}(\mathcal{T}, \mathcal{P})$ and for all $u \in \text{Obj}(\exists Y.B)$, that

$$C \models C([u]_{\sim}) \text{ iff } B \models C(u).$$

Since $\exists Y.B$ is a saturated repair, this implies that $\exists Z.C$ is saturated w.r.t. \mathcal{T} and does not entail (w.r.t. \mathcal{T}) any of the unwanted consequences specified by \mathcal{P} . The filtration $\exists Z.C$ need not be saturated w.r.t. \mathcal{R} , but we can show that its saturation w.r.t. \mathcal{R} does not entail additional instance relationships for atoms in $\text{Atoms}(\mathcal{T}, \mathcal{P}) \cup \text{Atoms}_{\mathcal{R}}(\mathcal{T}, \mathcal{P})$. This implies that, also w.r.t. $(\mathcal{T}, \mathcal{R})$, the filtration does not entail any of the unwanted consequences in \mathcal{P} . Finally, it is easy to check that $u \mapsto [u]_{\sim}$ is a homomorphism from the repair $\exists Y.B$ to the filtration $\exists Z.C$, and that $[u]_{\sim} \mapsto h(u)$ is a homomorphism from $\exists Z.C$ to the input qABox $\exists X.A$. Thus, the filtration $\exists Z.C$ is a repair with at most $m \cdot 2^n$ objects that entails $\exists Y.B$.

Since, for a fixed signature and up to renaming of variables, there are only finitely many qABoxes containing at most $m \cdot 2^n$ objects, we can effectively construct the set of optimal repairs of $\exists X.A$ for \mathcal{P} w.r.t. $(\mathcal{T}, \mathcal{R})$ by enumerating these qABoxes, then removing the ones that are not repairs, and finally removing from the remaining set the elements that are strictly entailed by an other element.

Theorem 12. *Let $\exists X.A$ be a qABox, $(\mathcal{T}, \mathcal{R})$ a terminating \mathcal{ELROI} terminology with regular RBox whose associated automata can effectively be computed, and \mathcal{P} an \mathcal{ELROI} repair request. Then the set of all optimal repairs of $\exists X.A$ for \mathcal{P} w.r.t. $(\mathcal{T}, \mathcal{R})$ can, up to equivalence, effectively be computed, and every repair is entailed by an optimal repair.*

The following example shows that the ‘‘automata atoms’’ in $\text{Atoms}_{\mathcal{R}}(\mathcal{T}, \mathcal{P})$ are needed for the filtration.

Example 13. Assume that $\mathcal{T} := \emptyset$, $\mathcal{R} := \{r \circ r \sqsubseteq s\}$, $\mathcal{A} := \{r(a, b), r(b, c), s(a, c)\}$, and $\mathcal{P} := \{\exists s.\top(a)\}$. The qABox $\exists \{x\}.\{r(a, b), r(x, c)\}$ is a (saturated) repair of $\exists \emptyset.A$ for \mathcal{P} w.r.t. $(\mathcal{T}, \mathcal{R})$. If we used only $\text{Atoms}(\mathcal{T}, \mathcal{P}) = \{\exists s.\top\}$ for the filtration, then the objects b and x would be identified since they behave the same w.r.t. this concept in the repair. Thus, $[a]_{\sim}$ would have $[x]_{\sim} = [b]_{\sim}$ as r -successor in the filtration, which in turn would have $[c]_{\sim}$ as r -successor. This shows that the filtration would have $\exists s.\top(a)$ as a consequence, and thus would not be a repair. The regular language $L_{\mathcal{R}}(s) = \{rr, s\}$ is accepted by a deterministic automaton with three states, q_0, q_1, q_2 , where q_0 is initial and q_2 is final, r -transitions from q_0 to q_1 and from q_1 to q_2 , and an s -transition from q_0 to q_2 . Since the object x belongs to $\exists q_1.\top$, but b does not, they are not identified in the filtration

that takes the concepts in $\text{Atoms}_{\mathcal{R}}(\mathcal{T}, \mathcal{P}) = \{\exists q_i. \top \mid 0 \leq i \leq 2\}$ into account.

Canonical Repairs Instead of blindly searching for optimal repairs among the very large set of “small” repairs, we now show how the considerably smaller set of canonical repairs, which contains all optimal repairs, can be constructed from repair seeds. Such a repair seed is of polynomial size, and it basically specifies which atoms in $\text{Atoms}(\mathcal{T}, \mathcal{P}) \cup \text{Atoms}_{\mathcal{R}}(\mathcal{T}, \mathcal{P})$ need to be removed for each individual.

From now on, we assume that $\exists Y.B$ is the saturation of $\exists X.A$ w.r.t. $(\mathcal{T}, \mathcal{R})$. Our canonical repairs will actually be computed from $\exists Y.B$. As mentioned in the introduction, to achieve optimality, it is not sufficient to remove assertions from this qABox. We must also generate anonymous copies of its objects. Basically, these copies are induced by pairs (u, \mathcal{K}) where u is an object in \mathcal{B} and $\mathcal{K} \subseteq \text{Atoms}_{\mathcal{R}}(\mathcal{T}, \mathcal{P})$ is a set of atoms C such that u is an instance of C in $\exists Y.B$. Putting an atom into \mathcal{K} means that the copy of u induced by (u, \mathcal{K}) should not be an instance of C .

Recall that $\text{Conj}(C)$ is the set of all top-level conjuncts of a concept description C . The set $\text{Conj}_{\mathcal{R}}(C)$ is obtained from $\text{Conj}(C)$ by replacing each existential restriction $\exists R.D \in \text{Conj}(C)$ with $\exists i_R.D$. The sets \mathcal{K} used to construct copies of u must be repair types for u .

Definition 14. Let u be an object name of $\exists Y.B$. A *repair type* for u is a set $\mathcal{K} \subseteq \text{Atoms}_{\mathcal{R}}(\mathcal{T}, \mathcal{P})$ satisfying:

- (RT1) If $C \in \mathcal{K}$, then $\mathcal{B} \models C(u)$.
- (RT2) If $D \in \text{Sub}(\mathcal{T}, \mathcal{P}) \cup \text{Atoms}_{\mathcal{R}}(\mathcal{T}, \mathcal{P})$ with $\mathcal{B} \models D(u)$ and $C \in \mathcal{K}$ with $D \sqsubseteq^{\mathcal{T}, \mathcal{R}} C$, then $\text{Conj}_{\mathcal{R}}(D) \cap \mathcal{K} \neq \emptyset$.
- (RT3) If $E \in \mathcal{P}_{\text{glo}}$ and $\mathcal{B} \models E(u)$, then $\text{Conj}_{\mathcal{R}}(E) \cap \mathcal{K} \neq \emptyset$.

The first condition says that only concept assertions that really hold for u need to be removed. The second condition ensures that concept assertions that are removed for u cannot be reintroduced by the terminology. The third condition has the effect that no copy can be an instance of a concept description that occurs in the global request.

In the canonical repairs, one of the copies of each individual will stand for this individual, whereas the other copies are variables. In addition, some individuals that are equal w.r.t. $\exists Y.B$ may no longer be equal in the repair. The repair seed makes these decisions explicit.

Definition 15. A *repair seed* \mathcal{S} consists of an equivalence relation $\approx_{\mathcal{S}}$ on $\text{Obj}(\exists Y.B)$ that is a refinement of $\approx_{\exists Y.B}$ (i.e., $\approx_{\mathcal{S}} \subseteq \approx_{\exists Y.B}$) and of a function that maps each equivalence class $[a]_{\mathcal{S}}$ of an individual a w.r.t. $\approx_{\mathcal{S}}$ to a repair type $\mathcal{S}_{[a]_{\mathcal{S}}}$ for a , such that the following conditions are fulfilled:

- (RS1) If $C(a) \in \mathcal{P}_{\text{loc}}$ and $\mathcal{B} \models C(a)$, then $\text{Conj}_{\mathcal{R}}(C) \cap \mathcal{S}_{[a]_{\mathcal{S}}} \neq \emptyset$.
- (RS2) If a, b are individuals and $\{a\} \in \text{Atoms}(\mathcal{T}, \mathcal{P})$, then $\{a\} \in \mathcal{S}_{[b]_{\mathcal{S}}}$ iff $a \approx_{\exists Y.B} b$ and $a \not\approx_{\mathcal{S}} b$.

The first condition guarantees that the repair induced by the seed satisfies the local request. The second condition ensures that the decision made by the seed that two individuals should no longer be equal is respected in the repair.

In contrast to the case of $\mathcal{E}\mathcal{L}$ (Baader et al. 2021), not every repair seed induces a canonical repair. The seed must

satisfy an additional admissibility conditions. Due to space constraints, we can neither provide the definition of admissibility nor the one of canonical repairs here. They can be found in (Baader and Kriegel 2022), as can the proof of the next theorem.

Theorem 16. For every admissible repair seed \mathcal{S} , the induced canonical repair can effectively be computed, is saturated w.r.t. $(\mathcal{T}, \mathcal{R})$, and is a repair of $\exists X.A$ for \mathcal{P} w.r.t. $(\mathcal{T}, \mathcal{R})$. Conversely, every repair of $\exists X.A$ for \mathcal{P} w.r.t. $(\mathcal{T}, \mathcal{R})$ is entailed by such a canonical repair.

Since all admissible repair seeds can effectively be generated, Theorem 12 can also be obtained as a corollary to this theorem. Even in the case without a terminology, not all canonical repairs need to be optimal (Baader et al. 2020), but we expect even the non-optimal ones to be quite good w.r.t. preserving consequences. One advantage of canonical repairs is that each one can be characterized by a polynomial-size repair seed, which can be generated by the knowledge engineers by making a polynomial number of decisions based on their domain knowledge. Another advantage is that the optimized approach for generating a canonical repair from a repair seed introduced in (Baader et al. 2021) for $\mathcal{E}\mathcal{L}$ can be extended to $\mathcal{E}\mathcal{L}\mathcal{R}\mathcal{O}\mathcal{I}$. More details on these advantages can be found in (Baader and Kriegel 2022).

4 Extensions and Applications

In this section we present several extensions to the repair framework. Section 4.1 shows how inconsistencies can be repaired that come into play when the bottom concept \perp is added. Section 4.2 deals with repair requests formulated as conjunctive queries. Finally, Section 4.3 briefly mentions additional extensions that cannot be presented in detail here.

4.1 Adding the Bottom Concept and Repairing Inconsistencies

The DL $\mathcal{E}\mathcal{L}\mathcal{R}\mathcal{O}\mathcal{I}(\perp)$ ($\mathcal{E}\mathcal{L}\mathcal{R}_{\text{reg}}\mathcal{O}\mathcal{I}(\perp)$) extends $\mathcal{E}\mathcal{L}\mathcal{R}\mathcal{O}\mathcal{I}$ ($\mathcal{E}\mathcal{L}\mathcal{R}_{\text{reg}}\mathcal{O}\mathcal{I}$) with the bottom concept \perp , which is always interpreted as the empty set. If \perp is available in the TBox, then qABoxes may become inconsistent w.r.t. terminologies. We call the quantified ABox $\exists X.A$ *consistent* w.r.t. a terminology $(\mathcal{T}, \mathcal{R})$ if there is a model of $\exists X.A$ and $(\mathcal{T}, \mathcal{R})$, and *inconsistent* otherwise. For instance, the qABox $\{A(a)\}$ is inconsistent w.r.t. the TBox $\{A \sqsubseteq \perp\}$.

Since any concept assertion (qABox) is entailed w.r.t. $(\mathcal{T}, \mathcal{R})$ by a qABox that is inconsistent w.r.t. $(\mathcal{T}, \mathcal{R})$, any non-empty repair request to an inconsistent qABox requires us also to get rid of the inconsistency. In addition, the definition of what is a repair needs to be revised since (Rep1) is trivially satisfied in case $\exists X.A$ is inconsistent w.r.t. $(\mathcal{T}, \mathcal{R})$. Any qABox $\exists Y.B$ satisfying (Rep2) and (Rep3) is thus a repair, even if $\exists Y.B$ is completely unrelated to $\exists X.A$. Hence, there cannot be an optimal repair since we can always extend a given repair by adding completely unrelated assertions.

Fortunately, in $\mathcal{E}\mathcal{L}\mathcal{R}\mathcal{O}\mathcal{I}(\perp)$ we can divide the TBox into a positive and an “unsatisfiable” part, where the unsatisfiable part plays a rôle when an inconsistency is derived, but has no effect otherwise. To be more precise, consider an $\mathcal{E}\mathcal{L}\mathcal{R}\mathcal{O}\mathcal{I}(\perp)$ TBox \mathcal{T} . Since each concept containing \perp is

equivalent to \perp , we can assume without loss of generality that each concept description occurring in \mathcal{T} is either \perp or does not contain \perp as a subconcept. After removing tautological CIs $\perp \sqsubseteq C$, it follows that \mathcal{T} is a disjoint union of a TBox \mathcal{T}_+ in which \perp does not occur (the *positive part*) and of a TBox \mathcal{T}_\perp containing only CIs of the form $C \sqsubseteq \perp$ where C does not contain \perp (the *unsatisfiable part*). We can characterize inconsistency by means of this partitioning of \mathcal{T} , and show that \mathcal{T}_\perp is only relevant for causing an inconsistency.

Proposition 17. *The following holds for every $\mathcal{ELROI}(\perp)$ terminology $(\mathcal{T}, \mathcal{R})$:*

1. *The quantified ABox $\exists X.A$ is inconsistent w.r.t. $(\mathcal{T}, \mathcal{R})$ iff there is a CI $C \sqsubseteq \perp$ in \mathcal{T}_\perp such that $\exists X.A \models_{\mathcal{T}_+, \mathcal{R}} \exists \{x\}. \{C(x)\}$.*
2. *If $\exists X.A$ is consistent w.r.t. $(\mathcal{T}, \mathcal{R})$, then $\exists X.A \models_{\mathcal{T}, \mathcal{R}} \exists Y.B$ iff $\exists X.A \models_{\mathcal{T}_+, \mathcal{R}} \exists Y.B$.*

Motivated by the second statement of this proposition, we now use \mathcal{T}_+ rather than \mathcal{T} in (Rep1), and of course additionally require the repair to be consistent. Also note that it does not make sense to use \perp in the repair request.

Definition 18. Consider a qABox $\exists X.A$, an \mathcal{ELROI} repair request \mathcal{P} , and an $\mathcal{ELROI}(\perp)$ terminology $(\mathcal{T}, \mathcal{R})$. An *inconsistency repair* of $\exists X.A$ for \mathcal{P} w.r.t. $(\mathcal{T}, \mathcal{R})$ is a qABox $\exists Y.B$ such that

- (IRep1) $\exists X.A \models_{\mathcal{T}_+, \mathcal{R}} \exists Y.B$
- (IRep2) $\exists Y.B$ is consistent w.r.t. $(\mathcal{T}, \mathcal{R})$,
- (IRep3) $\exists Y.B \not\models_{\mathcal{T}, \mathcal{R}} C(a)$ for each $C(a) \in \mathcal{P}_{\text{loc}}$, and
- (IRep4) $\exists Y.B \not\models_{\mathcal{T}, \mathcal{R}} \exists \{x\}. \{D(x)\}$ for each $D \in \mathcal{P}_{\text{glo}}$.

This inconsistency repair is *optimal* if it is not strictly entailed by another inconsistency repair w.r.t. $(\mathcal{T}, \mathcal{R})$.

Due to the second statement in Proposition 17, the notion of an inconsistency repair coincides with that of a repair as introduced in Definition 8 if $\exists X.A$ is consistent w.r.t. $(\mathcal{T}, \mathcal{R})$. If $\exists X.A$ is inconsistent w.r.t. $(\mathcal{T}, \mathcal{R})$, then the first statement in Proposition 17 shows that (IRep2) can be enforced by extending the global request with the concepts C for which $C \sqsubseteq \perp \in \mathcal{T}_\perp$. Given a repair request \mathcal{P} , we denote the extended request obtained this way as $\mathcal{P}^{\mathcal{T}_\perp}$.

Theorem 19. *Consider a qABox $\exists X.A$, an \mathcal{ELROI} repair request \mathcal{P} , and an $\mathcal{ELROI}(\perp)$ terminology $(\mathcal{T}, \mathcal{R})$. If $(\mathcal{T}, \mathcal{R})$ is inconsistent, then there are no inconsistency repairs of $\exists X.A$ w.r.t. $(\mathcal{T}, \mathcal{R})$. Otherwise, the (optimal) inconsistency repairs of $\exists X.A$ for \mathcal{P} w.r.t. $(\mathcal{T}, \mathcal{R})$ coincide with the (optimal) repairs of $\exists X.A$ for $\mathcal{P}^{\mathcal{T}_\perp}$ w.r.t. $(\mathcal{T}_+, \mathcal{R})$.*

If \mathcal{R} is regular and $(\mathcal{T}_+, \mathcal{R})$ is terminating, then we can apply the approach described in the previous section to compute all optimal inconsistency repairs.

4.2 Repairs for Conjunctive Queries

Until now, we have only allowed the use of \mathcal{ELROI} concept queries in the repair request. We now extend this to conjunctive queries (CQs). More precisely, we employ Boolean conjunctive queries (BCQs), i.e., CQs without answer variables. This is in line with the fact that we only considered

concept queries where the answer variable was either instantiated with an individual or existentially quantified. In (Grau and Kostylev 2019), CQs with answer variables are employed in the policy (which corresponds to our repair request), with the meaning that such a CQ should not have any answer tuple in the repair. This can clearly be expressed using the finitely many BCQs obtained by instantiating the answer variables with all answer tuples. As already mentioned above, BCQs and qABoxes are merely syntactic variants of each other (Baader et al. 2020). For this reason, we avoid introducing BCQs formally and use qABoxes instead.

Definition 20. A *qABox repair request* \mathcal{P} is a finite set of qABoxes. Given a qABox $\exists X.A$, a terminology $(\mathcal{T}, \mathcal{R})$, and a qABox repair request \mathcal{P} , a *repair* of $\exists X.A$ for \mathcal{P} w.r.t. $(\mathcal{T}, \mathcal{R})$ is a qABox $\exists Y.B$ that satisfies

- (CQRep1) $\exists X.A \models_{\mathcal{T}, \mathcal{R}} \exists Y.B$, and
- (CQRep2) $\exists Y.B \not\models_{\mathcal{T}, \mathcal{R}} \exists Z.C$ for each $\exists Z.C \in \mathcal{P}$.

It is *optimal* if it is not strictly entailed by another repair.

Since both concept assertions $C(a)$ and global repair requests $\exists \{x\}. \{C(x)\}$ for \mathcal{ELROI} concept descriptions C can be rewritten into equivalent qABoxes, using the first three rules in Figure 1, the repair requests and repairs introduced in Definition 8 are a special case of the qABox repair requests and repairs introduced here. We will now investigate under what conditions a rewriting in the other direction is possible.

Definition 21. An \mathcal{ELROI} *rewriting* of the qABox $\exists Z.C$ is an \mathcal{ELROI} concept description C such that $\exists Z.C$ and $\exists \{x\}. \{C(x)\}$ are equivalent.

By adapting the notion of c-acyclicity introduced in (Alexe et al. 2011), we can give (effectively checkable) conditions characterizing the existence of such a rewriting. Basically, the qABox is translated into an appropriate undirected graph, and the condition for c-acyclicity says that every cycle must contain an individual (see (Baader and Kriegel 2022) and (Alexe et al. 2011) for more details). Furthermore, we need the notion of a *core*, which is a qABox such that each endomorphism on it is bijective. Each qABox $\exists X.A$ has a computable and (up to renaming) unique core to which it is equivalent (Hell and Nešetřil 1992). It will be denoted in the following as $\text{core}(\exists X.A)$.

Proposition 22. *A qABox has an \mathcal{ELROI} rewriting iff its core is connected and c-acyclic.*

A detailed proof of this proposition can be found in (Baader and Kriegel 2022). The proof of the if-direction is constructive in the sense that it shows how the rewriting can be computed. Thus, if all qABoxes in a given qABox repair request are \mathcal{ELROI} rewritable, then we can reduce qABox repair to \mathcal{ELROI} repair.

Example 23. As an example of a qABox that is not c-acyclic we consider $\exists \{x, y\}. \{r(a, x), s(x, y), s(y, x)\}$. It has a cycle from x to y and then back that does not involve an individual. It is not \mathcal{ELROI} rewritable since an \mathcal{ELROI} concept could only enforce going back from y to the predecessor x if one of them were an individual whose name can be used in the concept. In contrast, the

$\text{qABox } \exists\{y\}.\{s(a, y), s(y, a)\}$, which is c-acyclic, has the \mathcal{ELROI} rewriting $\{a\} \sqcap \exists s.\exists s.\{a\}$. Note that the $\text{qABox } \exists\{x, y\}.\{s(x, y), r(x, y)\}$ is also not c-acyclic. Again, an \mathcal{ELROI} concept cannot enforce that there is a joint s - and r -successor of x . The $\text{qABox } \exists\{y\}.\{s(a, y), r(a, y)\}$ has the \mathcal{ELROI} rewriting $\exists r^-. \{a\} \sqcap \exists s^-. \{a\}$.

Considering Proposition 22, one might think that non-connectedness of $\text{core}(\exists Z.C)$ for $\exists Z.C \in \mathcal{P}$ could be an impediment to reducing qABox repair to \mathcal{ELROI} repair. However, this is not the case: it is sufficient that all connected components of $\text{core}(\exists Z.C)$ are \mathcal{ELROI} rewritable. To be more precise, let $\mathfrak{H}(\mathcal{P})$ be the set of all hitting sets of $\{\text{CoCo}(\text{core}(\exists Z.C)) \mid \exists Z.C \in \mathcal{P} \text{ and } \exists X.A \models^{\mathcal{T}, \mathcal{R}} \exists Z.C\}$, where the operator CoCo yields the set of connected components of an input qABox .

Lemma 24. $\exists Y.B$ is a repair of $\exists X.A$ for \mathcal{P} w.r.t. $(\mathcal{T}, \mathcal{R})$ iff there is a hitting set \mathcal{H} in $\mathfrak{H}(\mathcal{P})$ such that $\exists Y.B$ is a repair of $\exists X.A$ for \mathcal{H} w.r.t. $(\mathcal{T}, \mathcal{R})$.

According to our previous considerations, we can compute the optimal repairs for such a hitting set \mathcal{H} if each component in \mathcal{H} has an \mathcal{ELROI} rewriting. The elements of \mathcal{H} are connected components of the cores of the elements of \mathcal{P} . Since such a core is c-acyclic iff all its connected components are so, it is thus sufficient to require that the cores of the elements of \mathcal{P} are c-acyclic. Under this assumption, the set of all optimal repairs of $\exists X.A$ for \mathcal{P} w.r.t. $(\mathcal{T}, \mathcal{R})$ is then obtained by collecting the optimal repairs of $\exists X.A$ for \mathcal{H} w.r.t. $(\mathcal{T}, \mathcal{R})$ for all hitting set \mathcal{H} in $\mathfrak{H}(\mathcal{P})$, and then removing elements from this set that are strictly entailed by other elements.

Theorem 25. Let $\exists X.A$ be a qABox , $(\mathcal{T}, \mathcal{R})$ a terminating terminology with a regular RBox whose associated automata can effectively be computed, and \mathcal{P} be a qABox repair request. Then the set of all optimal repairs of $\exists X.A$ for \mathcal{P} w.r.t. $(\mathcal{T}, \mathcal{R})$ can be effectively computed if $\text{core}(\exists Z.C)$ is c-acyclic for all $\text{qABoxes } \exists Z.C$ in \mathcal{P} . In addition, each repair is then entailed by an optimal repair.

Without restrictions on the qABoxes in the repair request, the set of optimal repairs need not cover all repairs in the sense stated in the theorem, even if the qABox to be repaired is an ABox and the terminology is empty. In fact, it follows from (Nešetřil and Tardif 2000, Corollary 3.5) that the $\text{ABox } \{r(a, a)\}$ has no optimal repair for the repair request consisting of the $\text{qABox } \exists\{x\}.\{r(x, x)\}$. But the empty ABox is a repair, which is thus not entailed by an optimal one.

4.3 Further Extensions

The repair framework developed in this paper can also be used to deal with regular path expressions in the repair request, Horn- \mathcal{ALCOI} TBoxes , and qABoxes that have a static part that must not be changed. The basic idea is to create an \mathcal{ELROI} terminology over an extended signature that is a conservative extension of the input terminology, and in which such extensions can be expressed. Our repair approach is then applied w.r.t. this terminology. However, the repairs obtained this way may still contain names not occurring in the original signature, and thus these additional symbols need to be removed from these repairs appropriately.

We illustrate this for the case of *regular path expressions*, which are regular expressions over the alphabet of all roles. In repair requests, the concepts may now contain such expressions in place of roles. The semantics is defined by interpreting union, concatenation, and Kleene-star in the regular expressions as union, composition, and reflexive-transitive closure of binary relations, and the empty word as the identity relation. For example, the concept assertion $(\exists r^*.\{b\})(a) \in \mathcal{P}$ then says that, in the repair, there should not be an (empty or non-empty) r -path from a to b . To express the regular expression r^* , we introduce a new role name $\lceil r^* \rceil$ and extend the RBox with the RIs $\varepsilon \sqsubseteq \lceil r^* \rceil$, $r \sqsubseteq \lceil r^* \rceil$, and $\lceil r^* \rceil \circ \lceil r^* \rceil \sqsubseteq \lceil r^* \rceil$. In the repair request, we now use $\lceil r^* \rceil$ in place of r^* . A repair computed for this modified request may still contain the new name $\lceil r^* \rceil$, but we can simply remove all assertions containing it to obtain a repair in the original signature.

If a part of the given qABox is known to be correct, one may want to keep this part static when repairing (i.e., the repair should still imply this static part). Since our TBoxes are static and concept and role assertions can be expressed using nominals, the idea is now to move the static part of the qABox to the TBox . However, to express assertions involving variables, the signature needs to be extended by adding these variables as individual names.

More details on how to deal with these two extensions and on how Horn- \mathcal{ALCOI} TBoxes can be expressed can be found in (Baader and Kriegel 2022).

5 Conclusion

We have shown that the approaches for computing optimal repairs developed in our previous work can be extended to a considerably more expressive DL, which covers most of the DL \mathcal{EL}^{++} underlying the OWL2EL profile, but also has inverse roles. Our main result is that, in this setting, optimal repairs can effectively be computed and cover all repairs in the sense that every repair is entailed by an optimal one. In addition, we have demonstrated that this repair approach can deal with several other interesting repair problems.

The paper actually provides two proofs of the main result, one based on showing a small repair property by filtration, and another one based on the construction of canonical repairs. We believe the second approach to be more useful in practice. In fact, when repairing a given quantified ABox w.r.t. an \mathcal{ELROI} terminology, first computing all optimal repairs and then expecting the knowledge engineer to choose an appropriate one among (potentially) exponentially many exponentially large optimal repairs does not appear to be a practically viable repair approach. Since our canonical repairs are determined by polynomially large repair seeds, such a repair can be chosen by making polynomially many decisions regarding certain instance relationships. Once a repair seed is chosen, the induced canonical repair is always exponentially large. However, by adapting the optimized repair approach of (Baader et al. 2021) to our more expressive language, we can obtain considerably smaller optimized repairs.

Acknowledgements

This work was partially supported by the AI competence center ScaDS.AI Dresden/Leipzig and the German Research Foundation (DFG) in Project 430150274 (“Repairing Description Logic Ontologies”) and SFB/TRR 248 (“Foundations of Perspicuous Software Systems”). We thank the anonymous reviewers for their helpful comments. In particular, one of the reviewers provided us with Example 10 and the idea that filtration can be used to show the small repair property.

References

- Alexe, B.; ten Cate, B.; Kolaitis, P. G.; and Tan, W. C. 2011. Characterizing schema mappings via data examples. *ACM Trans. Database Syst.* 36(4):23:1–23:48.
- Baader, F., and Kriegel, F. 2022. Pushing optimal ABox repair from \mathcal{EL} towards more expressive Horn-DLs (extended version). LTCS-Report 22-02, Chair of Automata Theory, Institute of Theoretical Computer Science, Technische Universität Dresden, Dresden, Germany. <https://doi.org/10.25368/2022.131>.
- Baader, F., and Suntisrivaraporn, B. 2008. 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)*.
- Baader, F.; Horrocks, I.; Lutz, C.; and Sattler, U. 2017. *An Introduction to Description Logic*. Cambridge University Press.
- Baader, F.; Kriegel, F.; Nuradiansyah, A.; and Peñaloza, R. 2018. Making repairs in description logics more gentle. In Thielscher, M.; Toni, F.; and Wolter, F., eds., *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018*, 319–328. AAAI Press.
- Baader, F.; Kriegel, F.; Nuradiansyah, A.; and Peñaloza, R. 2020. Computing compliant anonymisations of quantified ABoxes w.r.t. \mathcal{EL} policies. In Pan, J. Z.; Tamma, V. A. M.; d’Amato, C.; Janowicz, K.; Fu, B.; Polleres, A.; Seneviratne, O.; and Kagal, L., eds., *The Semantic Web – ISWC 2020 – 19th International Semantic Web Conference, Proceedings, Part I*, volume 12506 of *Lecture Notes in Computer Science*, 3–20. Springer.
- Baader, F.; Koopmann, P.; Kriegel, F.; and Nuradiansyah, A. 2021. Computing optimal repairs of quantified ABoxes w.r.t. static \mathcal{EL} TBoxes. In Platzer, A., and Sutcliffe, G., eds., *Automated Deduction - CADE 28 - 28th International Conference on Automated Deduction, Proceedings*, volume 12699 of *Lecture Notes in Computer Science*, 309–326. Springer.
- Baader, F.; Borgwardt, S.; and Morawska, B. 2012. Extending unification in \mathcal{EL} towards general TBoxes. In Brewka, G.; Eiter, T.; and McIlraith, S. A., eds., *Principles of Knowledge Representation and Reasoning: Proceedings of the Thirteenth International Conference, KR 2012*. AAAI Press.
- Baader, F.; Brandt, S.; and Lutz, C. 2005. Pushing the \mathcal{EL} envelope. In Kaelbling, L. P., and Saffiotti, A., eds., *IJCAI-05, Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*, 364–369. Professional Book Center.
- Baader, F.; Lutz, C.; and Brandt, S. 2008. Pushing the \mathcal{EL} envelope further. In Clark, K., and Patel-Schneider, P. F., eds., *Proceedings of the Fourth OWLED Workshop on OWL: Experiences and Directions*, volume 496 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Calvanese, D.; Giacomo, G. D.; Lembo, D.; Lenzerini, M.; and Rosati, R. 2007. Tractable reasoning and efficient query answering in description logics: The *DL-Lite* family. *J. Autom. Reason.* 39(3):385–429.
- Calvanese, D.; Eiter, T.; and Ortiz, M. 2009. Regular path queries in expressive description logics with nominals. In Boutilier, C., ed., *IJCAI 2009, Proceedings of the 21st International Joint Conference on Artificial Intelligence*, 714–720.
- Du, J.; Qi, G.; and Fu, X. 2014. A practical fine-grained approach to resolving incoherent OWL 2 DL terminologies. In *Proc. of the 23rd ACM Int. Conf. on Information and Knowledge Management, (CIKM’14)*, 919–928.
- Grau, B. C., and Kostylev, E. V. 2019. Logical foundations of linked data anonymisation. *J. Artif. Intell. Res.* 64:253–314.
- Grau, B. C.; Horrocks, I.; Krötzsch, M.; Kupke, C.; Magka, D.; Motik, B.; and Wang, Z. 2013. Acyclicity notions for existential rules and their application to query answering in ontologies. *J. Artif. Intell. Res.* 47:741–808.
- Hell, P., and Nešetřil, J. 1992. The core of a graph. *Discret. Math.* 109(1-3):117–126.
- Horridge, M.; Parsia, B.; and Sattler, U. 2008. Laconic and precise justifications in OWL. In Sheth, A. P.; Staab, S.; Dean, M.; Paolucci, M.; Maynard, D.; Finin, T. W.; and Thirunarayan, K., eds., *The Semantic Web - ISWC 2008, 7th International Semantic Web Conference, Proceedings*, volume 5318 of *Lecture Notes in Computer Science*, 323–338. Springer.
- Horrocks, I.; Kutz, O.; and Sattler, U. 2006. The even more irresistible *SROIQ*. In Doherty, P.; Mylopoulos, J.; and Welty, C. A., eds., *Proceedings, Tenth International Conference on Principles of Knowledge Representation and Reasoning*, 57–67. AAAI Press.
- Kazakov, Y. 2008. *RIQ* and *SROIQ* are harder than *SHOIQ*. In Brewka, G., and Lang, J., eds., *Principles of Knowledge Representation and Reasoning: Proceedings of the Eleventh International Conference, KR 2008*, 274–284. AAAI Press.
- Kazakov, Y. 2009. Consequence-driven reasoning for Horn *SHIQ* ontologies. In Boutilier, C., ed., *IJCAI 2009, Proceedings of the 21st International Joint Conference on Artificial Intelligence*, 2040–2045.
- Kazakov, Y. 2010. An extension of complex role inclusion axioms in the description logic *SROIQ*. In Giesl, J., and Hähnle, R., eds., *Automated Reasoning, 5th International Joint Conference, IJCAR 2010, Proceedings*, volume 6173 of *Lecture Notes in Computer Science*, 472–486. Springer.

- Lam, J. S. C.; Sleeman, D. H.; Pan, J. Z.; and Vasconcelos, W. W. 2008. A fine-grained approach to resolving unsatisfiable ontologies. *J. Data Semant.* 10:62–95.
- Maurer, H. A.; Salomaa, A.; and Wood, D. 1980. Pure grammars. *Inf. Control.* 44(1):47–72.
- Nešetřil, J., and Tardif, C. 2000. Duality theorems for finite structures (characterising gaps and good characterisations). *J. Comb. Theory, Ser. B* 80(1):80–97.
- Ortiz, M.; Rudolph, S.; and Šimkus, M. 2010. Worst-case optimal reasoning for the Horn-DL fragments of OWL 1 and 2. In Lin, F.; Sattler, U.; and Truszczyński, M., eds., *Principles of Knowledge Representation and Reasoning: Proceedings of the Twelfth International Conference, KR 2010*.
- Parsia, B.; Sirin, E.; and Kalyanpur, A. 2005. Debugging OWL ontologies. In Ellis, A., and Hagino, T., eds., *Proc. of the 14th International Conference on World Wide Web (WWW'05)*, 633–640. ACM.
- Schlobach, S.; Huang, Z.; Cornet, R.; and Harmelen, F. 2007. Debugging incoherent terminologies. *J. Automated Reasoning* 39(3):317–349.
- Troquard, N.; Confalonieri, R.; Galliani, P.; Peñaloza, R.; Porello, D.; and Kutz, O. 2018. Repairing ontologies via axiom weakening. In McIlraith, S. A., and Weinberger, K. Q., eds., *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18)*, 1981–1988. AAAI Press.