

Optimal Fixed-Premise Repairs of \mathcal{EL} TBoxes

Francesco Kriegel 

Theoretical Computer Science, Technische Universität Dresden, Dresden, Germany
francesco.kriegel@tu-dresden.de

Abstract. Reasoners can be used to derive implicit consequences from an ontology. Sometimes unwanted consequences are revealed, indicating errors or privacy-sensitive information, and the ontology needs to be appropriately repaired. The classical approach is to remove just enough axioms such that the unwanted consequences vanish. However, this is often too rough since mere axiom deletion also erases many other consequences that might actually be desired. The goal should not be to remove a minimal number of axioms but to modify the ontology such that only a minimal number of consequences is removed, including the unwanted ones. Specifically, a repair should rather be logically entailed by the input ontology, instead of being a subset. To this end, we introduce a framework for computing fixed-premise repairs of \mathcal{EL} TBoxes. In the first variant the conclusions must be generalizations of those in the input TBox, while in the second variant no such restriction is imposed. In both variants, every repair is entailed by an optimal one and, up to equivalence, the set of all optimal repairs can be computed in exponential time. A prototypical implementation is provided. In addition, we show new complexity results regarding gentle repairs.

Keywords: Description logic · Optimal repair · TBox repair · Generalized-conclusion repair · Fixed-premise repair

1 Introduction

Description Logics (DLs) [4] are logic-based languages with model-theoretic semantics that are designed for knowledge representation and reasoning. Several DLs are fragments of first-order logic, but with restricted expressivity such that reasoning problems usually remain decidable. Knowledge represented as a DL ontology consists of a terminological part (the schema, TBox) and an assertional part (the data, ABox). The TBox expresses global knowledge on the underlying domain of interest, such as implicative rules and integrity constraints, and the ABox expresses local knowledge, such as assignment of objects to classes or relations between objects. DLs differ in their expressivity and there is always a trade-off to complexity of reasoning. Many reasoning tasks in lightweight DLs such as \mathcal{EL} [3] and $DL-Lite$ [12] are in P and thus tractable, but are N2EXP-complete in the very expressive DL $SROIQ$ [16,18], which is the logical foundation of the OWL2 Web Ontology Language.¹ However, the latter is a worst-case complexity, and efficient reasoning techniques [34] can often avoid reaching it.

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

Reasoners can be used to derive implicit consequences from an ontology. Sometimes unwanted consequences are revealed, indicating errors or privacy-sensitive information, and the ontology needs to be appropriately repaired. The classical approach is to remove just enough axioms such that the unwanted consequences vanish [14,29]. In particular, optimal classical repairs can be obtained by means of axiom pinpointing [10,11,31,32]: firstly, one determines all minimal subsets of the given ontology that entail the unwanted consequences (so-called *justifications*), secondly, one constructs a minimal set that contains at least one axiom from each justification (a so-called *hitting set*) and, thirdly, one removes from the erroneous ontology all axioms in the hitting set. In a similar way, inconsistency or incoherence of ontologies can be resolved — a task also called ontology debugging [17,22,30,33]. Proof visualizations can be used to guide the process of ontology repair [1], and it can be distributed and parallelized by means of decomposition [26]. Furthermore, there are connections to belief revision [13].

The classical repair approach is often too rough since mere axiom deletion also erases too many other consequences that might actually be desired. The goal should not be to remove a minimal number of axioms but to modify the ontology such that only a minimal number of consequences is removed, including the unwanted ones. Alternative repair techniques that are less dependent on the syntax should therefore be designed. To this end, a repair need not be a subset of the input ontology anymore, but must only be logically entailed by it.

A framework for constructing *gentle repairs* based on axiom weakening was developed [8]. The main difference to the classical repair approach is that, instead of being removed completely, one axiom from each justification is replaced by a logically weaker one such that the unwanted consequences cannot be derived anymore. The framework can be applied to every monotonic logic, and one only needs to devise a suitable weakening relation on axioms.² In terms of belief revision, gentle repairs correspond to pseudo-contractions [27].

In the DL \mathcal{EL} [3], concept descriptions are built from concept names and role names by conjunction and existential restriction, and a TBox is a finite set of concept inclusions (CIs), which are axioms of the form $C \sqsubseteq D$ where the premise C and the conclusion D are concept descriptions. For instance, the CI `MountainBike \sqsubseteq \exists hasPart.SuspensionFork \sqcap \exists isSuitableFor.OffRoadCycling` expresses that every mountain bike has a suspension fork and is suitable for off-road cycling. Such axioms can be weakened by specializing the premise or by generalizing the conclusion. Two weakening relations \succ^{syn} and \succ^{sub} for \mathcal{EL} CIs were devised [8], which instantiate the gentle repair framework for \mathcal{EL} TBoxes.

Repairs of \mathcal{EL} TBoxes can also be obtained by axiomatizing the logical intersection of the input TBox and the theory of a countermodel to the unwanted consequences [15], e.g., by means of the framework for axiomatizing \mathcal{EL} closure operators [19]. Such a countermodel can either be manually specified by the knowledge engineer or be automatically obtained by transforming a canonical model of the TBox, e.g., with the methods for repairing quantified ABoxes [9].

² There is always the trivial weakening relation that replaces each axiom with a tautology, for which each gentle repair is a classical repair.

The axiomatization method is very precise since it can introduce new premises in the resulting repair if necessary [15, Example 18]. From a theoretical perspective, this is a clear advantage simply because thereby a large amount of knowledge can be retained in the repair. From a practical perspective, however, this can be seen as a disadvantage as the resulting repairs might get considerably larger than the input TBox. In order to prevent such an increase in size, I have further proposed to construct a repair from a countermodel \mathcal{J} in a slightly different manner [15]: namely one keeps all premises unchanged and only generalizes the conclusions by means of \mathcal{J} , which yields an approach very close to the gentle repairs for the weakening relation \succ^{sub} .

The goal of this article is to elaborate the latter idea in detail. We introduce a framework for computing *generalized-conclusion repairs* of \mathcal{EL} TBoxes, where the premises must not be changed and the conclusions can be generalized. We first devise a canonical construction of such repairs from polynomial-size seeds, and then show that each generalized-conclusion repair is entailed by an optimal one and that, up to equivalence, the set of all optimal generalized-conclusion repairs can be computed in exponential time.

As an example, consider the TBox consisting of the single concept inclusion $\text{Bike} \sqsubseteq \exists \text{hasPart.SuspensionFork} \sqcap \exists \text{isSuitableFor.OffRoadCycling}$, which differs from the above in that the premise is replaced by Bike . It entails the false CIs $\text{Bike} \sqsubseteq \exists \text{hasPart.SuspensionFork}$ and $\text{Bike} \sqsubseteq \exists \text{isSuitableFor.OffRoadCycling}$. The (unique) optimal generalized-conclusion repair consists of the single CI $\text{Bike} \sqsubseteq \exists \text{hasPart.} \top \sqcap \exists \text{isSuitableFor.} \top$. In contrast, the classical repair approach deletes the single CI completely, yielding an empty repair, which only entails tautologies but does not entail that every bike has a part and is suitable for something.

In addition to developing the framework of generalized-conclusion repairs, we introduce *fixed-premise repairs*. The difference to the generalized-conclusion repairs is that the conclusions of CIs need not be generalizations anymore; only the premises must remain the same and the input TBox must entail each CI in the repair. Thereby even more consequences can be retained. Employing the same seeds as before, we show that every fixed-premise repair is entailed by an optimal one and that the set of all optimal fixed-premise repairs can be computed in exponential time.

Clearly, the above generalized-conclusion repair is not satisfactory if additional knowledge would be expressed in the given TBox, such as $\text{SuspensionFork} \sqsubseteq \text{Fork}$ and $\text{OffRoadCycling} \sqsubseteq \text{Cycling}$. Both additional CIs are obviously true in real world and should thus be retained in an optimal repair. Taking this into account, the (unique) optimal fixed-premise repair additionally contains the CI $\text{Bike} \sqsubseteq \exists \text{hasPart.Fork} \sqcap \exists \text{isSuitableFor.Cycling}$, and it preserves more consequences than the above generalized-conclusion repair, e.g., that every bike is suitable for cycling.

An experimental implementation is available.³ In addition, we provide new complexity results regarding gentle repairs w.r.t. the weakening relation \succ^{sub} . Due to space constraints, proofs can only be found in the extended version [20].

³ <https://github.com/francesco-kriegel/right-repairs-of-el-tboxes>

2 Preliminaries

Fix a *signature* Σ , which is a disjoint union of a set Σ_C of *concept names* and a set Σ_R of *role names*. In \mathcal{EL} , *concept descriptions* are inductively constructed by means of the grammar rule $C ::= \top \mid A \mid C \sqcap C \mid \exists r.C$ where A ranges over Σ_C and r over Σ_R . A *concept inclusion (CI)* is of the form $C \sqsubseteq D$ for concept descriptions C and D , where we call C the *premise* and D the *conclusion*. A *terminological box (TBox)* \mathcal{T} is a finite set of concept inclusions. The set of all premises in \mathcal{T} is denoted by $\text{Prem}(\mathcal{T})$.

The semantics is defined via models. An *interpretation* \mathcal{I} consists of a *domain* $\text{Dom}(\mathcal{I})$, which is a non-empty set, and an *interpretation function* $\cdot^{\mathcal{I}}$ that maps each concept name A to a subset $A^{\mathcal{I}}$ of $\text{Dom}(\mathcal{I})$ and that maps each role name r to a binary relation $r^{\mathcal{I}}$ over $\text{Dom}(\mathcal{I})$. The interpretation function is extended to all concept descriptions in the following recursive manner: $\top^{\mathcal{I}} := \text{Dom}(\mathcal{I})$, $(C \sqcap D)^{\mathcal{I}} := C^{\mathcal{I}} \cap D^{\mathcal{I}}$, and $(\exists r.C)^{\mathcal{I}} := \{x \mid (x, y) \in r^{\mathcal{I}} \text{ for some } y \in C^{\mathcal{I}}\}$. Furthermore, \mathcal{I} *satisfies* a CI $C \sqsubseteq D$ if $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$, written $\mathcal{I} \models C \sqsubseteq D$, and \mathcal{I} is a *model* of a TBox \mathcal{T} if it satisfies all CIs in \mathcal{T} , written $\mathcal{I} \models \mathcal{T}$. We say that \mathcal{T} *entails* $C \sqsubseteq D$ if $C \sqsubseteq D$ is satisfied in every model of \mathcal{T} , denoted as $\mathcal{T} \models C \sqsubseteq D$. We then also say that C *is subsumed by* D w.r.t. \mathcal{T} and write $C \sqsubseteq^{\mathcal{T}} D$. Subsumption in \mathcal{EL} can be decided in polynomial time [3]. With $C \sqsupset^{\mathcal{T}} D$ we abbreviate $C \sqsubseteq^{\mathcal{T}} D$ and $D \not\sqsubseteq^{\mathcal{T}} C$. Given sets \mathcal{K} and \mathcal{L} of \mathcal{EL} concept descriptions, we say that \mathcal{K} *is covered by* \mathcal{L} w.r.t. \mathcal{T} and write $\mathcal{K} \leq^{\mathcal{T}} \mathcal{L}$ if, for each $K \in \mathcal{K}$, there is some $L \in \mathcal{L}$ such that $K \sqsubseteq^{\mathcal{T}} L$.

An *atom* is either a concept name or an existential restriction $\exists r.C$. Order and repetitions of atoms in conjunctions as well as nestings of conjunctions are irrelevant. In this sense, each concept description C is a conjunction of atoms, which we call the *top-level conjuncts* of C , and the set of these is denoted by $\text{Conj}(C)$. Furthermore, we sometimes write $\prod\{C_1, \dots, C_n\}$ for $C_1 \sqcap \dots \sqcap C_n$. The (unique) *reduced form* C^r of a concept description C is obtained by exhaustively removing occurrences of atoms that subsume (w.r.t. \emptyset) another atom in the same conjunction. C is equivalent to C^r , and two concept descriptions are equivalent iff they have the same reduced form [21]. The subsumption order \sqsubseteq^{\emptyset} restricted to reduced concept descriptions is a partial order and not just a pre-order [9].

We denote by $\text{Sub}(\alpha)$ the set of all concept descriptions that occur as sub-concepts in α , and $\text{Atoms}(\alpha)$ is the set of atoms occurring in α . Given a set \mathcal{K} of atoms, $\text{Max}(\mathcal{K})$ denotes the subset consisting of all \sqsubseteq^{\emptyset} -maximal atoms, i.e., $\text{Max}(\mathcal{K}) := \{K \mid K \in \mathcal{K} \text{ and there is no } K' \in \mathcal{K} \text{ such that } K \sqsubset^{\emptyset} K'\}$. If all atoms in \mathcal{K} are reduced, then $\text{Max}(\mathcal{K})$ does not contain \sqsubseteq^{\emptyset} -comparable atoms.

Let \mathcal{I} be an interpretation and X a subset of $\text{Dom}(\mathcal{I})$. A *most specific concept description (MSC)* of X w.r.t. \mathcal{I} is a concept description C that satisfies $X \subseteq C^{\mathcal{I}}$ and, for each concept description D , $X \subseteq D^{\mathcal{I}}$ implies $C \sqsubseteq^{\emptyset} D$. The MSC of X w.r.t. \mathcal{I} is unique up to equivalence and is denoted as $X^{\mathcal{I}}$. Due to cycles in the interpretation, MSCs might not be expressible in \mathcal{EL} , but MSCs always exist in an extension of \mathcal{EL} with greatest fixed-points, e.g., in \mathcal{EL}_{si} [23]. The latter DL extends \mathcal{EL} with *simulation quantifiers* $\exists^{\text{sim}}(\mathcal{I}, x)$ where the semantics of such concept descriptions is defined by: $y \in (\exists^{\text{sim}}(\mathcal{I}, x))^{\mathcal{I}}$ if there is a

simulation from \mathcal{I} to \mathcal{J} that contains (x, y) . As shown in [19, Proposition 4.1.6], the MSC $X^{\mathcal{I}}$ is equivalent to $\exists^{\text{sim}}(\wp(\mathcal{I}), X)$, where the *powering* $\wp(\mathcal{I})$ has domain $\text{Dom}(\wp(\mathcal{I})) := \wp(\text{Dom}(\mathcal{I}))$, and $A^{\wp(\mathcal{I})}$ consists of all subsets X such that $X \subseteq A^{\mathcal{I}}$, and $r^{\wp(\mathcal{I})}$ consists of all pairs (X, Y) such that Y is a minimal hitting set of $\{\{y \mid (x, y) \in r^{\mathcal{I}}\} \mid x \in X\}$. A CI $C \sqsubseteq D$ is satisfied in \mathcal{I} iff $C^{\mathcal{I}\mathcal{I}} \sqsubseteq^{\emptyset} D$, and $X \subseteq C^{\mathcal{I}}$ is equivalent to $X^{\mathcal{I}} \sqsubseteq^{\emptyset} C$ for each subset $X \subseteq \text{Dom}(\mathcal{I})$ and for each \mathcal{EL}_{si} concept description C .

A *least common subsumer (LCS)* of concept descriptions C and D is a concept description E such that $C \sqsubseteq^{\emptyset} E$ as well as $D \sqsubseteq^{\emptyset} E$ and, for each concept description F , $C \sqsubseteq^{\emptyset} F$ and $D \sqsubseteq^{\emptyset} F$ implies $E \sqsubseteq^{\emptyset} F$. The LCS of C and D is unique up to equivalence and we denote it by $C \vee D$. It can be computed as the product of the graphs representing C and D . In particular, the LCS of an \mathcal{EL} concept description C and an \mathcal{EL}_{si} concept description $\exists^{\text{sim}}(\mathcal{I}, x)$ is always expressible in \mathcal{EL} and the following recursion allows us to construct it:

$$C \vee \exists^{\text{sim}}(\mathcal{I}, x) \equiv^{\emptyset} \prod \{ A \mid A \in \text{Conj}(C) \text{ and } x \in A^{\mathcal{I}} \} \\ \sqcap \prod \{ \exists r. (D \vee \exists^{\text{sim}}(\mathcal{I}, y)) \mid \exists r. D \in \text{Conj}(C) \text{ and } (x, y) \in r^{\mathcal{I}} \}.$$

Furthermore, the MSC $X^{\mathcal{I}}$ is equivalent to the LCS of all $\exists^{\text{sim}}(\mathcal{I}, x)$ where $x \in X$.

3 Generalized-Conclusion Repairs of \mathcal{EL} TBoxes

In this section we develop the framework for computing generalized-conclusion repairs of \mathcal{EL} TBoxes. We begin with defining basic notions.

Definition 1. Let \mathcal{T} and \mathcal{U} be \mathcal{EL} TBoxes. We say that \mathcal{U} is a generalized-conclusion weakening (GC-weakening) of \mathcal{T} , written $\mathcal{T} \succeq_{\text{GC}} \mathcal{U}$ if, for each CI $C \sqsubseteq D$ in \mathcal{U} , there is a CI $E \sqsubseteq F$ in \mathcal{T} such that $C = E$ and $F \sqsubseteq^{\emptyset} D$.

GC-weakening is strictly stronger than entailment, i.e., $\mathcal{T} \succeq_{\text{GC}} \mathcal{U}$ implies $\mathcal{T} \models \mathcal{U}$ but the converse need not hold. For instance, $\{A \sqcap B \sqsubseteq \exists r. (A \sqcap B), C \sqsubseteq A \sqcap \exists r. A\}$ has the GC-weakening $\{A \sqcap B \sqsubseteq \exists r. A \sqcap \exists r. B, C \sqsubseteq \exists r. A\}$, and it entails $\{A \sqcap B \sqsubseteq \exists r. (A \sqcap \exists r. A)\}$, which is not a GC-weakening.

Definition 2. A repair request \mathcal{P} is a finite set of \mathcal{EL} concept inclusions. A TBox \mathcal{T} complies with \mathcal{P} if it does not entail any CI in \mathcal{P} , i.e., it holds that $\mathcal{T} \not\models C \sqsubseteq D$ for each $C \sqsubseteq D \in \mathcal{P}$. A countermodel to \mathcal{P} is an interpretation in which none of the CIs in \mathcal{P} is satisfied.

Definition 3. Given an \mathcal{EL} TBox \mathcal{T} and a repair request \mathcal{P} , a generalized-conclusion repair (GC-repair) of \mathcal{T} for \mathcal{P} is an \mathcal{EL} TBox \mathcal{U} that is a GC-weakening of \mathcal{T} and complies with \mathcal{P} . We further call \mathcal{U} optimal if there is no other GC-repair \mathcal{V} such that $\mathcal{V} \succeq_{\text{GC}} \mathcal{U}$ but $\mathcal{U} \not\succeq_{\text{GC}} \mathcal{V}$.

Throughout the whole section we assume that \mathcal{T} is an \mathcal{EL} TBox and that \mathcal{P} is a repair request, and the goal is to construct a generalized-conclusion repair (preferably an optimal one). Of course, if \mathcal{P} contains a tautology, then no repair exists. We therefore assume that this is not the case. Without loss of generality, all concept descriptions in \mathcal{T} and \mathcal{P} must be reduced.

Induced Countermodels. In the first step, we transform a canonical model of the input TBox \mathcal{T} into countermodels to \mathcal{P} , which are used in the next section to devise a canonical construction of generalized-conclusion repairs. The construction of each countermodel is guided by a repair seed.

Definition 4. A repair seed is a TBox \mathcal{S} that complies with \mathcal{P} and consists of CIs of the form $C \sqsubseteq F$ for a premise $C \in \text{Prem}(\mathcal{T})$ and an atom $F \in \text{Atoms}(\mathcal{P}, \mathcal{T})$ where $C \sqsubseteq^{\mathcal{T}} F$.

The completion algorithm for \mathcal{EL} is a decision procedure for the subsumption problem (and also for the instance problem). In the correctness proof a canonical model of the TBox is constructed that involves all subconcepts occurring in the TBox [3]. While this algorithm works in a rule-based manner, thus implicitly constructing the canonical model step by step, there is also a closed-form representation [25]. Resembling the latter we define the *canonical model* \mathcal{I} with domain $\text{Dom}(\mathcal{I}) := \{x_C \mid C \in \text{Sub}(\mathcal{P}, \mathcal{T})\}$ and its interpretation function is given by $A^{\mathcal{I}} := \{x_C \mid C \sqsubseteq^{\mathcal{T}} A\}$ for each $A \in \Sigma_C$ and $r^{\mathcal{I}} := \{(x_C, x_D) \mid C \sqsubseteq^{\mathcal{T}} \exists r.D\}$ for each $r \in \Sigma_R$. Then \mathcal{I} is a model of \mathcal{T} , and $x_C \in E^{\mathcal{I}}$ iff $C \sqsubseteq^{\mathcal{T}} E$ for each subconcept $C \in \text{Sub}(\mathcal{P}, \mathcal{T})$ and for each \mathcal{EL} concept description E [20].

The transformation of the canonical model \mathcal{I} is based on modification types. These describe how copies of objects in the domain of \mathcal{I} are modified in order to create objects of a countermodel.

Definition 5. Let $x_C \in \text{Dom}(\mathcal{I})$. A modification type for x_C is a subset \mathcal{K} of $\text{Atoms}(\mathcal{P}, \mathcal{T})$ where $x_C \in K^{\mathcal{I}}$ for each $K \in \mathcal{K}$, and $K_1 \not\sqsubseteq^{\mathcal{I}} K_2$ for each two $K_1, K_2 \in \mathcal{K}$. Given a repair seed \mathcal{S} , we say that \mathcal{K} respects \mathcal{S} if additionally $\{D\} \leq^{\mathcal{S}} \mathcal{K}$ implies $\{D\} \leq^{\mathcal{I}} \mathcal{K}$ for each $D \in \text{Sub}(\mathcal{P}, \mathcal{T})$ where $x_C \in D^{\mathcal{I}}$.

Each repair seed \mathcal{S} induces a countermodel to \mathcal{P} . Its domain consists of all copies of objects in the canonical model \mathcal{I} that are annotated with an \mathcal{S} -respecting modification type. The definition of the interpretation function guarantees that each such copy does not satisfy any atom in the modification type.

Definition 6. Let \mathcal{S} be a repair seed. The induced countermodel $\mathcal{J}_{\mathcal{S}}$ has the domain $\text{Dom}(\mathcal{J}_{\mathcal{S}})$ consisting of all objects $x_{C, \mathcal{K}}$ where $x_C \in \text{Dom}(\mathcal{I})$ and \mathcal{K} is a modification type for x_C that respects \mathcal{S} , and its interpretation function is defined by $A^{\mathcal{J}_{\mathcal{S}}} := \{x_{C, \mathcal{K}} \mid x_C \in A^{\mathcal{I}} \text{ and } A \notin \mathcal{K}\}$ for each concept name $A \in \Sigma_C$ and $r^{\mathcal{J}_{\mathcal{S}}} := \{(x_{C, \mathcal{K}}, x_{D, \mathcal{L}}) \mid (x_C, x_D) \in r^{\mathcal{I}} \text{ and } \text{Succ}(\mathcal{K}, r, x_D) \leq^{\mathcal{I}} \mathcal{L}\}$ for each role name $r \in \Sigma_R$, where $\text{Succ}(\mathcal{K}, r, x_D) := \{E \mid \exists r.E \in \mathcal{K} \text{ and } x_D \in E^{\mathcal{I}}\}$.

We can show that an object $x_{C, \mathcal{K}}$ satisfies an \mathcal{EL} concept description E in $\mathcal{J}_{\mathcal{S}}$ iff x_C satisfies E in \mathcal{I} and \mathcal{K} does not contain an atom subsuming E [20]. Now consider an unwanted CI $C \sqsubseteq D$ in the repair request \mathcal{P} . Since \mathcal{S} complies with \mathcal{P} , there is a top-level conjunct D' in D such that $\mathcal{S} \not\models C \sqsubseteq D'$. We can thus construct an \mathcal{S} -respecting modification type \mathcal{K} for x_C that contains an atom subsuming D' but none subsuming C . It follows that the copy $x_{C, \mathcal{K}}$ satisfies the premise C but not the conclusion D , i.e., $\mathcal{J}_{\mathcal{S}}$ is indeed a countermodel to $C \sqsubseteq D$.

Proposition 7. For each repair seed \mathcal{S} , the induced countermodel $\mathcal{J}_{\mathcal{S}}$ is a countermodel to \mathcal{P} .

Canonical Generalized-Conclusion Repairs. Next, we show how each repair seed \mathcal{S} induces a GC-repair. We obtain it by generalizing each conclusion according to countermodel $\mathcal{J}_{\mathcal{S}}$, namely we take each concept inclusion $C \sqsubseteq D$ in the given TBox \mathcal{T} and replace D with the least common subsumer of D and the most specific concept description E for which the CI $C \sqsubseteq E$ is satisfied in $\mathcal{J}_{\mathcal{S}}$.

Definition 8. *Each repair seed \mathcal{S} induces the TBox*

$$\text{rep}_{\text{GC}}(\mathcal{T}, \mathcal{S}) := \{ C \sqsubseteq D \vee C^{\mathcal{J}_{\mathcal{S}}\mathcal{J}_{\mathcal{S}}} \mid C \sqsubseteq D \in \mathcal{T} \}.$$

The following lemma shows that $\text{rep}_{\text{GC}}(\mathcal{T}, \mathcal{S})$ has exactly those TBoxes as GC-weakenings that are GC-weakenings of \mathcal{T} and of which $\mathcal{J}_{\mathcal{S}}$ is a model.

Lemma 9. *$\text{rep}_{\text{GC}}(\mathcal{T}, \mathcal{S}) \succeq_{\text{GC}} \mathcal{U}$ iff $\mathcal{T} \succeq_{\text{GC}} \mathcal{U}$ and $\mathcal{J}_{\mathcal{S}} \models \mathcal{U}$*

As $\text{rep}_{\text{GC}}(\mathcal{T}, \mathcal{S})$ is a GC-weakening of itself, we infer that $\mathcal{J}_{\mathcal{S}}$ is a model of $\text{rep}_{\text{GC}}(\mathcal{T}, \mathcal{S})$. According to Proposition 7, $\mathcal{J}_{\mathcal{S}}$ is a countermodel to \mathcal{P} , and so $\text{rep}_{\text{GC}}(\mathcal{T}, \mathcal{S})$ complies with \mathcal{P} . It is further easy to see that $\text{rep}_{\text{GC}}(\mathcal{T}, \mathcal{S})$ is a GC-weakening of \mathcal{T} . We have thus shown that the following holds.

Proposition 10. *If \mathcal{S} is a repair seed, then $\text{rep}_{\text{GC}}(\mathcal{T}, \mathcal{S})$ is a GC-repair.*

If the repair request \mathcal{P} does not contain a tautological CI, then the empty set is already a repair seed, i.e., $\text{rep}_{\text{GC}}(\mathcal{T}, \emptyset)$ is a GC-repair of \mathcal{T} for \mathcal{P} . Furthermore, the induced GC-repairs are complete in the sense that every GC-repair is a GC-weakening of $\text{rep}_{\text{GC}}(\mathcal{T}, \mathcal{S})$ for some repair seed \mathcal{S} .

Proposition 11. *If \mathcal{U} is a GC-repair of \mathcal{T} for \mathcal{P} , then there is a repair seed \mathcal{S} such that $\text{rep}_{\text{GC}}(\mathcal{T}, \mathcal{S}) \succeq_{\text{GC}} \mathcal{U}$.*

Proof Sketch. Given a GC-repair \mathcal{U} , a repair seed $\mathcal{S}_{\mathcal{U}}^*$ is obtained as the least fixed point of the equation $\mathcal{S} = \{ C \sqsubseteq F \mid C \sqsubseteq D' \in \mathcal{U}, F \in \text{Atoms}(\mathcal{P}, \mathcal{T}), \text{ and } D' \sqsubseteq^{\mathcal{S}} F \}$. It has the important property that $\{D'\} \leq^{\mathcal{S}_{\mathcal{U}}^*} \mathcal{K}$ implies $\{C\} \leq^{\mathcal{S}_{\mathcal{U}}^*} \mathcal{K}$ for each CI $C \sqsubseteq D' \in \mathcal{U}$ and for each modification type \mathcal{K} . With this property we can easily show that $x_{E, \mathcal{K}} \in C^{\mathcal{J}_{\mathcal{S}_{\mathcal{U}}^*}}$ implies $x_{E, \mathcal{K}} \in (D')^{\mathcal{J}_{\mathcal{S}_{\mathcal{U}}^*}}$ for each CI $C \sqsubseteq D' \in \mathcal{U}$, and thus the induced countermodel $\mathcal{J}_{\mathcal{S}_{\mathcal{U}}^*}$ is a model of \mathcal{U} . Lemma 9 yields that \mathcal{U} is a GC-weakening of $\text{rep}_{\text{GC}}(\mathcal{T}, \mathcal{S}_{\mathcal{U}}^*)$. \square

Each repair seed is of polynomial size, and there are at most exponentially many seeds. Even with a naïve approach, we can compute all seeds in exponential time and thus also all induced GC-repairs. Then we must filter out the non-optimal ones, e.g., by comparing each two repairs w.r.t. \succeq_{GC} . Each comparison needs polynomial time [3], and we obtain the following main result.

Theorem 12. *The set of all optimal GC-repairs of an \mathcal{EL} TBox \mathcal{T} for a repair request \mathcal{P} can be computed in exponential time, and each GC-repair is a GC-weakening of an optimal one.*

In the below example, an optimal GC-repair is not polynomial-time computable.

Example 13. For the repair request $\{\exists r.A \sqsubseteq \exists r.B\}$, the TBox $\{\exists r.A \sqsubseteq \exists r.(P_1 \sqcap Q_1 \sqcap \dots \sqcap P_n \sqcap Q_n), P_1 \sqcap Q_1 \sqsubseteq B, \dots, P_n \sqcap Q_n \sqsubseteq B\}$ has the optimal GC-repair $\{\exists r.A \sqsubseteq \prod\{\exists r.(X_1 \sqcap \dots \sqcap X_n) \mid X_i \in \{P_i, Q_i\} \text{ for each } i \in \{1, \dots, n\}\}, P_1 \sqcap Q_1 \sqsubseteq B, \dots, P_n \sqcap Q_n \sqsubseteq B\}$. It has exponential size.

Computing a Canonical Generalized-Conclusion Repair. In the last step, we are concerned with the question how the GC-repair induced by a seed \mathcal{S} can efficiently be computed. Recall that, as explained in the preliminaries, each conclusion $D \vee C^{\mathcal{J}_S \mathcal{J}_S}$ can be obtained as the product of the \mathcal{EL} concept description D and the \mathcal{EL}_{si} concept description $\exists^{\text{sim}}(\wp(\mathcal{J}_S), C^{\mathcal{J}_S})$, or alternatively as the product of D and all $\exists^{\text{sim}}(\mathcal{J}_S, x_{E, \mathcal{K}})$ where $x_{E, \mathcal{K}} \in C^{\mathcal{J}_S}$. However, computing the induced GC-repair $\text{rep}_{\text{GC}}(\mathcal{T}, \mathcal{S})$ in this way is very inefficient since \mathcal{J}_S has exponential size.

The first important observation is that the concept description $C^{\mathcal{J}_S \mathcal{J}_S}$ is already equivalent to $\exists^{\text{sim}}(\mathcal{J}_S, x_{C, \mathcal{S}[C]})$ where $\mathcal{S}[C]$ is the largest modification type for x_C that respects \mathcal{S} and does not contain an atom subsuming C . This follows from the fact that there is a simulation on \mathcal{J}_S that contains the pair $(x_{C, \mathcal{S}[C]}, x_{E, \mathcal{K}})$ for each object $x_{E, \mathcal{K}}$ in the extension $C^{\mathcal{J}_S}$. Secondly, in order to compute the LCS $D \vee \exists^{\text{sim}}(\mathcal{J}_S, x_{C, \mathcal{S}[C]})$ it is not necessary to start from $x_{C, \mathcal{S}[C]}$ in the product construction, but it suffices to start from $x_{D, \mathcal{S}[C \sqsubseteq D]}$ where $\mathcal{S}[C \sqsubseteq D]$ is the largest modification type for x_D that respects \mathcal{S} and does not contain an atom subsuming C . Thirdly, when computing the product of D and $\exists^{\text{sim}}(\mathcal{J}_S, x_{D, \mathcal{S}[C \sqsubseteq D]})$ we do not need to consider all objects $x_{E, \mathcal{K}}$ that are reachable from $x_{D, \mathcal{S}[C \sqsubseteq D]}$ in \mathcal{J}_S , but only those where E is a filler of an existential restriction that occurs in D . As main result we obtain the following proposition.

Definition 14. Given a subconcept $E \in \text{Sub}(\mathcal{P}, \mathcal{T})$ and a modification type \mathcal{K} for x_E that respects \mathcal{S} , we define the restriction $E \upharpoonright_{\mathcal{K}}$ by the following recursion.

$$E \upharpoonright_{\mathcal{K}} := \prod \{ A \mid A \in \text{Conj}(E) \text{ and } A \notin \mathcal{K} \} \\ \prod \prod \left\{ \exists r. F \upharpoonright_{\mathcal{L}} \mid \begin{array}{l} \exists r. F \in \text{Conj}(E), \text{ and } \mathcal{L} \text{ is a } \leq^0\text{-minimal mod. type} \\ \text{for } x_F \text{ that respects } \mathcal{S} \text{ and where } \text{Succ}(\mathcal{K}, r, x_F) \leq^0 \mathcal{L} \end{array} \right\}$$

Proposition 15. Given a repair seed \mathcal{S} , it holds that $D \vee C^{\mathcal{J}_S \mathcal{J}_S} \equiv^0 D \upharpoonright_{\mathcal{S}[C \sqsubseteq D]}$ for each CI $C \sqsubseteq D$ in \mathcal{T} , and thus the induced GC-repair $\text{rep}_{\text{GC}}(\mathcal{T}, \mathcal{S})$ is equivalent to the TBox $\{ C \sqsubseteq D \upharpoonright_{\mathcal{S}[C \sqsubseteq D]} \mid C \sqsubseteq D \in \mathcal{T} \}$, where the modification type $\mathcal{S}[C \sqsubseteq D]$ is defined as $\text{Max}\{ K \mid K \in \text{Atoms}(\mathcal{P}, \mathcal{T}), C \not\sqsubseteq^S K, \text{ and } D \sqsubseteq^T K \}$.

Two Observations. The below example illustrates that entailment between repair seeds need not imply entailment between the induced GC-repairs.

Example 16. For the TBox $\mathcal{T} := \{ A \sqsubseteq B, C \sqsubseteq \exists r.(A \sqcap B) \}$ and the repair request $\mathcal{P} := \{ C \sqsubseteq \exists r.B \}$, there are two optimal GC-repairs: $\mathcal{U}_1 := \{ A \sqsubseteq B, C \sqsubseteq \exists r. \top \}$, induced by the seed $\mathcal{S}_1 := \{ A \sqsubseteq B \}$, and $\mathcal{U}_2 := \{ A \sqsubseteq \top, C \sqsubseteq \exists r.A \}$, induced by $\mathcal{S}_2 := \emptyset$. Now, \mathcal{U}_1 does not entail \mathcal{U}_2 , although \mathcal{S}_1 entails \mathcal{S}_2 .

The next example shows that, possibly contradicting intuition, it does not suffice that a repair seed consists only of CIs $C \sqsubseteq F$ where $C \sqsubseteq D \in \mathcal{T}$ and $F \in \text{Atoms}(\mathcal{P}, \mathcal{T})$ such that $D \sqsubseteq^0 F$. We definitely sometimes need CIs $C \sqsubseteq F$ where $C \sqsubseteq^T F$, as per Definition 4. Notably, the only optimal repair in the following example can be described by the latter CIs.

Example 17. Consider the TBox $\mathcal{T} := \{A \sqsubseteq \exists r.\exists r.(B \sqcap C), \exists r.B \sqsubseteq B\}$ and the repair request $\mathcal{P} := \{A \sqsubseteq \exists r.\exists r.C\}$. The unique optimal GC-repair is $\{A \sqsubseteq \exists r.\exists r.B, \exists r.B \sqsubseteq B\}$. It is induced only by the seeds $\{A \sqsubseteq \exists r.B, \exists r.B \sqsubseteq B\}$ and $\{A \sqsubseteq B, A \sqsubseteq \exists r.B, \exists r.B \sqsubseteq B\}$. Specifically the seed CI $A \sqsubseteq \exists r.B$ would not be allowed if we simplified the definition of a seed as explained above.

Another GC-repair is $\{A \sqsubseteq \exists r.\exists r.B, \exists r.B \sqsubseteq \top\}$, which is induced by the empty seed \emptyset , but also by $\{A \sqsubseteq B\}$, $\{A \sqsubseteq \exists r.B\}$, and $\{A \sqsubseteq B, A \sqsubseteq \exists r.B\}$.

The above example also shows that a repair need not entail its seed, and that a repair can be induced by multiple seeds. Conducted experiments support the claim that each GC-repair might be induced by a unique seed with minimal cardinality and such that every CI in the seed is also entailed by the repair.

4 Fixed-Premise Repairs of \mathcal{EL} TBoxes

We have seen in the introduction that simply generalizing the conclusions of the input TBox \mathcal{T} might not yield satisfactory repairs. Therefore, we will now construct repairs that can retain more consequences. It is still required that each premise in the repair is also a premise in \mathcal{T} , but apart from that we do not impose further conditions except that the repair must, of course, be entailed by \mathcal{T} .

Definition 18. Consider TBoxes \mathcal{T} and \mathcal{U} . We say that \mathcal{T} fixed-premise entails (FP-entails) \mathcal{U} , written $\mathcal{T} \models_{\text{FP}} \mathcal{U}$, if $\text{Prem}(\mathcal{T}) = \text{Prem}(\mathcal{U})$ and $\mathcal{T} \models \mathcal{U}$.

$\mathcal{T} \succeq_{\text{GC}} \mathcal{U}$ implies $\mathcal{T} \models_{\text{FP}} \mathcal{U}$ and the latter implies $\mathcal{T} \models \mathcal{U}$, but the converse implications need not hold. This means that the relation \models_{FP} is between \succeq_{GC} and \models . Thus, repairs based on this new relation are, usually, better than GC-repairs.

Definition 19. Let \mathcal{T} be an \mathcal{EL} TBox and \mathcal{P} a repair request. A fixed-premise repair (FP-repair) of \mathcal{T} for \mathcal{P} is an \mathcal{EL} TBox \mathcal{U} that is FP-entailed by \mathcal{T} and complies with \mathcal{P} . We further call \mathcal{U} optimal if there is no other FP-repair \mathcal{V} such that $\mathcal{V} \models_{\text{FP}} \mathcal{U}$ and $\mathcal{U} \not\models_{\text{FP}} \mathcal{V}$.

Obviously, each GC-repair is an FP-repair but the converse does not hold.

By reusing the notion of a repair seed as well as the results on GC-repairs in Section 3, we obtain the following characterization of (optimal) FP-repairs. First of all, each repair seed \mathcal{S} induces an FP-repair: we take each CI $C \sqsubseteq D$ in the input TBox \mathcal{T} and replace the conclusion D with the most specific concept description E for which the CI $C \sqsubseteq E$ is satisfied in the induced countermodel $\mathcal{J}_{\mathcal{S}}$. Note that now D is not generalized anymore by computing an LCS.

Definition 20. Each repair seed \mathcal{S} induces the TBox

$$\text{rep}_{\text{FP}}(\mathcal{T}, \mathcal{S}) := \{C \sqsubseteq C^{\mathcal{J}_{\mathcal{S}}\mathcal{J}_{\mathcal{S}}} \mid C \in \text{Prem}(\mathcal{T})\}.$$

Recall that each conclusion $C^{\mathcal{J}_{\mathcal{S}}\mathcal{J}_{\mathcal{S}}}$ is equivalent to the \mathcal{EL}_{si} concept description $\exists^{\text{sim}}(\mathcal{J}_{\mathcal{S}}, x_{C, \mathcal{S}[C]})$, where $\mathcal{S}[C]$ is the largest modification type for x_C that

respects \mathcal{S} and does not cover $\{C\}$, i.e., $\mathcal{S}[C] := \text{Max}\{K \mid K \in \text{Atoms}(\mathcal{P}, \mathcal{T}), C \sqsubseteq^{\mathcal{S}} K, \text{ and } C \sqsubseteq^{\mathcal{T}} K\}$. Analogously to the GC-repairs, every TBox $\text{rep}_{\text{FP}}(\mathcal{T}, \mathcal{S})$ is an FP-repair and each FP-repair is FP-entailed by $\text{rep}_{\text{FP}}(\mathcal{T}, \mathcal{S})$ for some repair seed \mathcal{S} .

Proposition 21. *For each repair seed \mathcal{S} , the TBox $\text{rep}_{\text{FP}}(\mathcal{T}, \mathcal{S})$ is an FP-repair.*

Proposition 22. *For each FP-repair \mathcal{U} of \mathcal{T} for \mathcal{P} , there is a repair seed \mathcal{S} such that $\text{rep}_{\text{FP}}(\mathcal{T}, \mathcal{S}) \models_{\text{FP}} \mathcal{U}$.*

We obtain the following main result of this section. Its proof is analogous to Theorem 12, but uses the argument that entailment between \mathcal{EL}_{si} TBoxes can be decided in polynomial time [23].

Theorem 23. *The set of all optimal FP-repairs of an \mathcal{EL} TBox \mathcal{T} for a repair request \mathcal{P} can be computed in exponential time, and each FP-repair is FP-entailed by an optimal one.*

We have seen in Example 17 that a repair seed might not be entailed by its induced GC-repair. This is *not* the case for its induced FP-repair.

Lemma 24. *Each repair seed \mathcal{S} is entailed by its induced FP-repair $\text{rep}_{\text{FP}}(\mathcal{T}, \mathcal{S})$.*

Contrary to the GC-repairs, not every FP-repair is an \mathcal{EL} TBox but might require cyclic \mathcal{EL}_{si} concept descriptions [23] as conclusions to be optimal. For instance, consider the TBox $\{A \sqsubseteq \exists r.A\}$ that is also the repair request. The unique optimal FP-repair consists of the single CI

$$A \sqsubseteq \exists^{\text{sim}}(\rightarrow \overset{A}{\circ} \xrightarrow{r} \circ \xrightarrow{r} \overset{A}{\circ} \rhd r).$$

If a standard \mathcal{EL} TBox is required as result, one might rewrite the repair by introducing fresh concept names (used as quantified monadic second-order variables). For the above optimal repair this yields the TBox $\exists\{X, Y, Z\}. \{A \sqsubseteq X, X \sqsubseteq A \sqcap \exists r.Y, Y \sqsubseteq \exists r.Z, Z \sqsubseteq A \sqcap \exists r.Z\}$. One could also try to compute a uniform interpolant [24, 28] of the latter in order to get rid of the additional symbols and so obtain a usual \mathcal{EL} TBox. Alternatively, one could unfold the cyclic conclusions into \mathcal{EL} concept descriptions up to a certain role-depth bound.

If the TBox \mathcal{T} is cycle-restricted [2], then the canonical model \mathcal{I} is acyclic and so is the induced countermodel $\mathcal{J}_{\mathcal{S}}$ for each repair seed \mathcal{S} . The FP-repair $\text{rep}_{\text{FP}}(\mathcal{T}, \mathcal{S})$ then only has acyclic \mathcal{EL}_{si} concept descriptions as conclusions and these can be rewritten into \mathcal{EL} concept descriptions.

5 Complexity of Maximally Strong \succ^{sub} -Weakenings

As mentioned in the introduction, a framework for computing *gentle repairs* based on axiom weakening was developed, and two weakening relations that operate on \mathcal{EL} CIs were introduced [8]. We briefly recall the *modified gentle*

repair algorithm. As input, fix an ontology \mathcal{O} that is partitioned into a static part \mathcal{O}_s and a refutable part \mathcal{O}_r , as well as an axiom α , the unwanted consequence, that follows from \mathcal{O} but not already from \mathcal{O}_s . A *repair* is an ontology \mathcal{O}' such that $\mathcal{O} \models \mathcal{O}'$ but $\mathcal{O}_s \cup \mathcal{O}' \not\models \alpha$. In order to obtain such a repair, we repeatedly compute a justification J for α and replace one axiom $\beta \in J$ by a weaker one.⁴ Specifically, a *justification* for α is a minimal subset $J \subseteq \mathcal{O}_r$ such that $\mathcal{O}_s \cup J \models \alpha$. After at most exponentially many iterations a repair has been obtained.

A *weakening relation* is a pre-order \succ on axioms such that $\beta \succ \gamma$ implies that γ is weaker than β . Such relations are used to guide the selection of a weaker axiom in the above iteration. Specifically, when processing a justification J for α and a selected axiom $\beta \in J$, we should replace β by a *maximally strong weakening*, which is an axiom γ such that $\beta \succ \gamma$ and $\mathcal{O}_s \cup (J \setminus \{\beta\}) \cup \{\gamma\} \not\models \alpha$, but $\mathcal{O}_s \cup (J \setminus \{\beta\}) \cup \{\delta\} \models \alpha$ for all δ where $\beta \succ \delta \succ \gamma$. This prevents the loss of too many other consequences (apart from α). However, maximally strong weakenings need not exist for every weakening relation.

The *syntactic weakening relation* \succ^{syn} on \mathcal{EL} CIs removes subconcepts from the conclusions. Maximally strong \succ^{syn} -weakenings always exist in all directions,⁵ all of them can be computed in exponential time, one can be computed in polynomial time, and recognizing them is **coNP**-complete.

The *semantic weakening relation* \succ^{sub} replaces conclusions of \mathcal{EL} CIs by more general concepts, i.e., $C \sqsubseteq D \succ^{\text{sub}} C' \sqsubseteq D'$ if $C = C'$, $D \sqsubset^\emptyset D'$, and $C' \sqsubseteq D' \not\models C \sqsubseteq D$. It has only been known that maximally strong \succ^{sub} -weakenings always exist in all directions,⁵ all of them can effectively be computed, and recognizing them is **coNP**-hard. As a side result from Section 3, we obtain the following.

Proposition 25. *If the unwanted consequence α is a CI, then all maximally strong \succ^{sub} -weakenings of an axiom β in a justification J for α can be computed in exponential time.*

The following modification of [8, Example 30] shows that a single maximally strong \succ^{sub} -weakening cannot always be computed in polynomial time.

Example 26. Take the ontology \mathcal{O} with $\mathcal{O}_s := \{P_i \sqcap Q_i \sqsubseteq B \mid i \in \{1, \dots, n\}\}$ and $\mathcal{O}_r := \{\beta\}$ for $\beta := \exists r.A \sqsubseteq \exists r.(P_1 \sqcap Q_1 \sqcap \dots \sqcap P_n \sqcap Q_n)$, and the unwanted consequence $\alpha := \exists r.A \sqsubseteq \exists r.B$. Then $J := \{\beta\}$ is a justification for α . There is exactly one maximally strong \succ^{sub} -weakening of β in J , namely $\exists r.A \sqsubseteq \prod\{\exists r.(X_1 \sqcap \dots \sqcap X_n) \mid X_i \in \{P_i, Q_i\} \text{ for each } i \in \{1, \dots, n\}\}$. Since this weakening has exponential size, it cannot be computed in polynomial time.

Finally, recognizing maximally strong \succ^{sub} -weakenings is also in **coNP**.

Proposition 27. *The problem of deciding whether an \mathcal{EL} CI γ is a maximally strong \succ^{sub} -weakening of an \mathcal{EL} CI β in a justification J for α is **coNP**-complete.*

⁴ We say that γ is *weaker* than β if β entails γ but γ does not entail β .

⁵ That is, each weakening of an axiom β in a justification J is weaker than a maximally strong weakening of β in J —where a *weakening* of β in J is an axiom γ such that $\beta \succ \gamma$ and $\mathcal{O}_s \cup (J \setminus \{\beta\}) \cup \{\gamma\} \not\models \alpha$.

6 Conclusion

We have introduced a framework for computing generalized-conclusion repairs of \mathcal{EL} TBoxes, where the premises must not be changed and the conclusions can be generalized. Up to equivalence, the set of all optimal generalized-conclusion repairs can be computed in exponential time. Each generalized-conclusion repair is entailed by an optimal one and, furthermore, each optimal generalized-conclusion repair can be described by a repair seed that has polynomial size. In addition, we have extended the framework to the fixed-premise repairs, with the difference that the conclusions need not be generalizations anymore. This usually leads to better repairs, but with the disadvantage that the conclusions in an optimal repair might be cyclic and can thus only be expressed in an extension of \mathcal{EL} with greatest fixed-point semantics or by introducing fresh concept names. Not affected by the latter, all optimal fixed-premise repairs can be computed in exponential time too, and each fixed-premise repair is entailed by an optimal one, which is induced by a polynomial-size repair seed. An experimental implementation is available, which interacts with the user to construct the seed from which the repair is built.

An interesting task for future research is to combine this approach to repairing TBoxes with the approach to repairing quantified ABoxes [5]. This should be possible by, firstly, adapting the notion of a repair seed such that it can additionally contain concept assertions and role assertions and, secondly, suitably adapting the transformation of the saturation/canonical model into a countermodel from which the final repair is constructed. Another interesting question is how the approach can be extended to more expressive DLs, such as \mathcal{EL} with the bottom concept \perp , nominals $\{a\}$, inverse roles r^- , and role inclusions $R_1 \circ \dots \circ R_n \sqsubseteq S$. Ideas from the latest extension of quantified ABox repairs to the DL $\mathcal{ELROI}(\perp)$ might be helpful [6, 7]. An extension with nominals would immediately add support for ABox axioms, since each concept assertion $C(a)$ is equivalent to the CI $\{a\} \sqsubseteq C$ and each role assertion is equivalent to $\{a\} \sqsubseteq \exists r.\{b\}$. Furthermore, it should not be hard to add support for a partitioning of the TBox into a static and a refutable part, or for a set of wanted consequences that must still be entailed by the repair. Also, it would be interesting to find a suitable partial order on repair seeds such that minimality of the seed is equivalent to optimality of the induced repair, similar to the qABox repairs [9]. Last, it would be interesting to investigate whether and how the quality of the repairs can be improved if also new premises can be introduced by the repair process. Currently, this can be done by manually extending the input TBox to be repaired.

Acknowledgements. The author thanks Franz Baader for his valuable comments that helped to improve this article. He further thanks the anonymous reviewers for their helpful feedback. The author is funded by Deutsche Forschungsgemeinschaft (DFG) in Project 430150274 (“Repairing Description Logic Ontologies”).

References

1. Alrabbaa, C., Baader, F., Dachsel, R., Flemisch, T., Koopmann, P.: Visualising proofs and the modular structure of ontologies to support ontology repair. In: Borgwardt, S., Meyer, T. (eds.) Proceedings of the 33rd International Workshop on Description Logics (DL 2020) co-located with the 17th International Conference on Principles of Knowledge Representation and Reasoning (KR 2020), Online Event [Rhodes, Greece], September 12th to 14th, 2020. CEUR Workshop Proceedings, vol. 2663. CEUR-WS.org (2020), <http://ceur-ws.org/Vol-2663/paper-2.pdf>
2. Baader, F., Borgwardt, S., Morawska, B.: A goal-oriented algorithm for unification in \mathcal{EL} w.r.t. cycle-restricted TBoxes. In: Kazakov, Y., Lembo, D., Wolter, F. (eds.) Proceedings of the 2012 International Workshop on Description Logics, DL-2012, Rome, Italy, June 7-10, 2012. CEUR Workshop Proceedings, vol. 846. CEUR-WS.org (2012), http://ceur-ws.org/Vol-846/paper_1.pdf
3. Baader, F., Brandt, S., Lutz, C.: Pushing the \mathcal{EL} envelope. In: Kaelbling, L.P., Saffiotti, A. (eds.) IJCAI-05, Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence, Edinburgh, Scotland, UK, July 30 - August 5, 2005. pp. 364–369. Professional Book Center (2005), <http://ijcai.org/Proceedings/05/Papers/0372.pdf>
4. Baader, F., Horrocks, I., Lutz, C., Sattler, U.: An Introduction to Description Logic. Cambridge University Press (2017), <https://doi.org/10.1017/9781139025355>
5. 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.) Automated Deduction - CADE 28 - 28th International Conference on Automated Deduction, Virtual Event, July 12-15, 2021, Proceedings. Lecture Notes in Computer Science, vol. 12699, pp. 309–326. Springer (2021), https://doi.org/10.1007/978-3-030-79876-5_18
6. Baader, F., Kriegel, F.: Pushing optimal ABox repair from \mathcal{EL} towards more expressive Horn-DLs. In: Kern-Isberner, G., Lakemeyer, G., Meyer, T. (eds.) Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning, KR 2022, Haifa, Israel, July 31 – August 5, 2022. pp. 22–32 (2022), <https://doi.org/10.24963/kr.2022/3>
7. Baader, F., Kriegel, F.: 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 (2022), <https://doi.org/10.25368/2022.131>
8. Baader, F., Kriegel, F., Nuradiansyah, A., Peñaloza, R.: Making repairs in description logics more gentle. In: Thielscher, M., Toni, F., Wolter, F. (eds.) 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>
9. Baader, F., Kriegel, F., Nuradiansyah, A., Peñaloza, R.: 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., Kagal, L. (eds.) The Semantic Web - ISWC 2020 - 19th International Semantic Web Conference, Athens, Greece, November 2-6, 2020, Proceedings, Part I. Lecture Notes in Computer Science, vol. 12506, pp. 3–20. Springer (2020), https://doi.org/10.1007/978-3-030-62419-4_1

10. Baader, F., Peñaloza, R.: Axiom pinpointing in general tableaux. *J. Log. Comput.* **20**(1), 5–34 (2010), <https://doi.org/10.1093/logcom/exn058>
11. Baader, F., Peñaloza, R., Suntisrivaraporn, B.: Pinpointing in the description logic \mathcal{EL} . In: Calvanese, D., Franconi, E., Haarslev, V., Lembo, D., Motik, B., Turhan, A., Tessaris, S. (eds.) Proceedings of the 2007 International Workshop on Description Logics (DL2007), Brixen-Bressanone, near Bozen-Bolzano, Italy, 8-10 June, 2007. CEUR Workshop Proceedings, vol. 250. CEUR-WS.org (2007), http://ceur-ws.org/Vol-250/paper_16.pdf
12. Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M., Rosati, R.: *DL-Lite*: tractable description logics for ontologies. In: Veloso, M.M., Kambhampati, S. (eds.) Proceedings, The Twentieth National Conference on Artificial Intelligence and the Seventeenth Innovative Applications of Artificial Intelligence Conference, July 9-13, 2005, Pittsburgh, Pennsylvania, USA. pp. 602–607. AAAI Press / The MIT Press (2005), <https://www.aaai.org/Papers/AAAI/2005/AAAI05-094.pdf>
13. C  be, R., Wassermann, R.: Ontology repair through partial meet contraction. In: Booth, R., Casini, G., Klarman, S., Richard, G., Varzinczak, I.J. (eds.) Proceedings of the International Workshop on Defeasible and Ampliative Reasoning, DARE 2015, co-located with the 24th International Joint Conference on Artificial Intelligence (IJCAI 2015), Buenos Aires, Argentina, July 27, 2015. CEUR Workshop Proceedings, vol. 1423. CEUR-WS.org (2015), http://ceur-ws.org/Vol-1423/DARE-15_2.pdf
14. Greiner, R., Smith, B.A., Wilkerson, R.W.: A correction to the algorithm in Reiter’s theory of diagnosis. *Artif. Intell.* **41**(1), 79–88 (1989), [https://doi.org/10.1016/0004-3702\(89\)90079-9](https://doi.org/10.1016/0004-3702(89)90079-9)
15. Hieke, W., Kriegel, F., Nuradiansyah, A.: Repairing \mathcal{EL} TBoxes by means of countermodels obtained by model transformation. In: Homola, M., Ryzhikov, V., Schmidt, R.A. (eds.) Proceedings of the 34th International Workshop on Description Logics (DL 2021) part of Bratislava Knowledge September (BAKS 2021), Bratislava, Slovakia, September 19th to 22nd, 2021. CEUR Workshop Proceedings, vol. 2954. CEUR-WS.org (2021), <http://ceur-ws.org/Vol-2954/paper-17.pdf>
16. Horrocks, I., Kutz, O., Sattler, U.: The even more irresistible *SR \mathcal{OIQ}* . In: Doherty, P., Mylopoulos, J., Welty, C.A. (eds.) Proceedings, Tenth International Conference on Principles of Knowledge Representation and Reasoning, Lake District of the United Kingdom, June 2-5, 2006. pp. 57–67. AAAI Press (2006), <https://www.aaai.org/Papers/KR/2006/KR06-009.pdf>
17. Kalyanpur, A., Parsia, B., Sirin, E., Cuenca Grau, B.: Repairing unsatisfiable concepts in OWL ontologies. In: Sure, Y., Domingue, J. (eds.) The Semantic Web: Research and Applications, 3rd European Semantic Web Conference, ESWC 2006, Budva, Montenegro, June 11-14, 2006, Proceedings. Lecture Notes in Computer Science, vol. 4011, pp. 170–184. Springer (2006), https://doi.org/10.1007/11762256_15
18. Kazakov, Y.: *RIQ* and *SR \mathcal{OIQ}* are harder than *SH \mathcal{OIQ}* . In: Brewka, G., Lang, J. (eds.) Principles of Knowledge Representation and Reasoning: Proceedings of the Eleventh International Conference, KR 2008, Sydney, Australia, September 16-19, 2008. pp. 274–284. AAAI Press (2008), <https://www.aaai.org/Papers/KR/2008/KR08-027.pdf>
19. Kriegel, F.: Constructing and Extending Description Logic Ontologies using Methods of Formal Concept Analysis. Doctoral thesis, Technische Universit  t Dresden, Dresden, Germany (2019), <http://nbn-resolving.de/urn:nbn:de:bsz:14-qucosa2-360998>

20. Kriegel, F.: Optimal fixed-premise repairs of \mathcal{EL} TBoxes (extended version). LTCS-Report 22-04, Chair of Automata Theory, Institute of Theoretical Computer Science, Technische Universität Dresden, Dresden, Germany (2022), <https://doi.org/10.25368/2022.321>
21. Küsters, R.: Non-Standard Inferences in Description Logics, Lecture Notes in Computer Science, vol. 2100. Springer (2001), <https://doi.org/10.1007/3-540-44613-3>
22. Lam, J.S.C., Pan, J.Z., Sleeman, D.H., Vasconcelos, W.W.: A fine-grained approach to resolving unsatisfiable ontologies. In: 2006 IEEE / WIC / ACM International Conference on Web Intelligence (WI 2006), 18-22 December 2006, Hong Kong, China. pp. 428–434. IEEE Computer Society (2006), <https://doi.org/10.1109/WI.2006.11>
23. Lutz, C., Piro, R., Wolter, F.: Enriching \mathcal{EL} -concepts with greatest fixpoints. In: Coelho, H., Studer, R., Wooldridge, M.J. (eds.) ECAI 2010 - 19th European Conference on Artificial Intelligence, Lisbon, Portugal, August 16-20, 2010, Proceedings. Frontiers in Artificial Intelligence and Applications, vol. 215, pp. 41–46. IOS Press (2010), <https://doi.org/10.3233/978-1-60750-606-5-41>
24. Lutz, C., Seylan, I., Wolter, F.: An automata-theoretic approach to uniform interpolation and approximation in the description logic \mathcal{EL} . In: Brewka, G., Eiter, T., McIlraith, S.A. (eds.) Principles of Knowledge Representation and Reasoning: Proceedings of the Thirteenth International Conference, KR 2012, Rome, Italy, June 10-14, 2012. AAAI Press (2012), <https://dl.acm.org/doi/abs/10.5555/3031843.3031877>
25. Lutz, C., Wolter, F.: Conservative extensions in the lightweight description logic \mathcal{EL} . In: Pfenning, F. (ed.) Automated Deduction - CADE-21, 21st International Conference on Automated Deduction, Bremen, Germany, July 17-20, 2007, Proceedings. Lecture Notes in Computer Science, vol. 4603, pp. 84–99. Springer (2007), https://doi.org/10.1007/978-3-540-73595-3_7
26. Ma, Y., Peñaloza, R.: Towards parallel repair: An ontology decomposition-based approach. In: Bienvenu, M., Ortiz, M., Rosati, R., Šimkus, M. (eds.) Informal Proceedings of the 27th International Workshop on Description Logics, Vienna, Austria, July 17-20, 2014. CEUR Workshop Proceedings, vol. 1193, pp. 633–645. CEUR-WS.org (2014), http://ceur-ws.org/Vol-1193/paper_61.pdf
27. Matos, V.B., Guimarães, R., David Santos, Y., Wassermann, R.: Pseudo-contractions as gentle repairs. In: Lutz, C., Sattler, U., Tinelli, C., Turhan, A., Wolter, F. (eds.) Description Logic, Theory Combination, and All That - Essays Dedicated to Franz Baader on the Occasion of His 60th Birthday. Lecture Notes in Computer Science, vol. 11560, pp. 385–403. Springer (2019), https://doi.org/10.1007/978-3-030-22102-7_18
28. Nikitina, N., Rudolph, S.: ExpExpExplosion: uniform interpolation in general \mathcal{EL} terminologies. In: De Raedt, L., Bessière, C., Dubois, D., Doherty, P., Frasconi, P., Heintz, F., Lucas, P.J.F. (eds.) ECAI 2012 - 20th European Conference on Artificial Intelligence. Including Prestigious Applications of Artificial Intelligence (PAIS-2012) System Demonstrations Track, Montpellier, France, August 27-31, 2012. Frontiers in Artificial Intelligence and Applications, vol. 242, pp. 618–623. IOS Press (2012), <https://doi.org/10.3233/978-1-61499-098-7-618>
29. Reiter, R.: A theory of diagnosis from first principles. *Artif. Intell.* **32**(1), 57–95 (1987), [https://doi.org/10.1016/0004-3702\(87\)90062-2](https://doi.org/10.1016/0004-3702(87)90062-2), see the erratum [14].
30. Scharrenbach, T., Grütter, R., Waldvogel, B., Bernstein, A.: Structure preserving TBox repair using defaults. In: Haarslev, V., Toman, D., Weddell, G.E. (eds.)

- Proceedings of the 23rd International Workshop on Description Logics (DL 2010), Waterloo, Ontario, Canada, May 4-7, 2010. CEUR Workshop Proceedings, vol. 573. CEUR-WS.org (2010), http://ceur-ws.org/Vol-573/paper_17.pdf
31. Schlobach, S.: Diagnosing terminologies. In: Veloso, M.M., Kambhampati, S. (eds.) Proceedings, The Twentieth National Conference on Artificial Intelligence and the Seventeenth Innovative Applications of Artificial Intelligence Conference, July 9-13, 2005, Pittsburgh, Pennsylvania, USA. pp. 670–675. AAAI Press / The MIT Press (2005), <https://www.aaai.org/Papers/AAAI/2005/AAAI05-105.pdf>
 32. Schlobach, S., Cornet, R.: Non-standard reasoning services for the debugging of description logic terminologies. In: Gottlob, G., Walsh, T. (eds.) IJCAI-03, Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence, Acapulco, Mexico, August 9-15, 2003. pp. 355–362. Morgan Kaufmann (2003), <http://ijcai.org/Proceedings/03/Papers/053.pdf>
 33. 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>
 34. Tena Cucala, D., Cuenca Grau, B., Horrocks, I.: Pay-as-you-go consequence-based reasoning for the description logic *SRCTLQ*. *Artif. Intell.* **298**, 103518 (2021), <https://doi.org/10.1016/j.artint.2021.103518>