Technische Universität Dresden
Institute for Theoretical Computer Science
Chair for Automata Theory

LTCS–Report

Optimal Fixed-Premise Repairs of $\mathcal{EL}$ TBoxes (Extended Version)

Francesco Kriegel

LTCS-Report 22-04

This is an extended version of an article accepted at the
45th German Conference on Artificial Intelligence (KI 2022).

Postal Address:
Lehrstuhl für Automatentheorie
Institut für Theoretische Informatik
TU Dresden
01062 Dresden

http://lat.inf.tu-dresden.de

Visiting Address:
Nöthnitzer Str. 46
Dresden

Optimal Fixed-Premise Repairs of $\mathcal{EL}$ TBoxes (Extended Version)

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* [13] are in P and thus tractable, but are $\mathsf{N2EXP}$-complete in the very expressive DL $\mathcal{SROIQ}$ [17,19], which is the logical foundation of the OWL 2 Web Ontology Language.[1] However, the latter is a worst-case complexity, and efficient reasoning techniques [36] 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 [15,31]. In particular, optimal classical repairs can be obtained by means of axiom pinpointing [11,12,33,34]: 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 [18,24,32,35]. Proof visualizations can be used to guide the process of ontology repair [1], and it can be distributed and parallelized by means of decomposition [28]. Furthermore, there are connections to belief revision [14].

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 [29].

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 $\mathsf{MountainBike} \sqsubseteq \exists \mathsf{hasPart}.\mathsf{SuspensionFork} \sqcap \exists \mathsf{isSuitableFor}.\mathsf{OffRoadCycling}$ 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 $\succ^{\mathsf{syn}}$ and $\succ^{\mathsf{sub}}$ 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 [16], e.g., by means of the framework for axiomatizing $\mathcal{EL}$ closure operators [20]. 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 [16, 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 [16]: 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 $\succ^{\mathsf{sub}}$.

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 $\mathsf{Bike} \sqsubseteq \exists \mathsf{hasPart}.\mathsf{SuspensionFork} \sqcap \exists \mathsf{isSuitableFor}.\mathsf{OffRoadCycling}$, which differs from the above in that the premise is replaced by Bike. It entails the false CIs $\mathsf{Bike} \sqsubseteq \exists \mathsf{hasPart}.\mathsf{SuspensionFork}$ and $\mathsf{Bike} \sqsubseteq \exists \mathsf{isSuitableFor}.\mathsf{OffRoadCycling}$. The (unique) optimal generalized-conclusion repair consists of the single CI $\mathsf{Bike} \sqsubseteq \exists \mathsf{hasPart}.\top \sqcap \exists \mathsf{isSuitableFor}.\top$. 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 $\mathsf{SuspensionFork} \sqsubseteq \mathsf{Fork}$ and $\mathsf{OffRoadCycling} \sqsubseteq \mathsf{Cycling}$. 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 $\mathsf{Bike} \sqsubseteq \exists \mathsf{hasPart}.\mathsf{Fork} \sqcap \exists \mathsf{isSuitableFor}.\mathsf{Cycling}$, 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 $\succ^{\mathsf{sub}}$.

[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 $\mathsf{Prem}(\mathcal{T})$.

The semantics is defined via models. An *interpretation* $\mathcal{I}$ consists of a *domain* $\mathsf{Dom}(\mathcal{I})$, 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 $\mathsf{Dom}(\mathcal{I})$ and that maps each role name r to a binary relation $r^{\mathcal{I}}$ over $\mathsf{Dom}(\mathcal{I})$. The interpretation function is extended to all concept descriptions in the following recursive manner: $\top^{\mathcal{I}} := \mathsf{Dom}(\mathcal{I})$, $(C \sqcap D)^{\mathcal{I}} := C^{\mathcal{I}} \cap D^{\mathcal{I}}$, 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 $C \sqsubset^{\mathcal{T}} D$ 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 $\mathsf{Conj}(C)$. Furthermore, we sometimes write $\bigsqcap\{C_1, \ldots, C_n\}$ for $C_1 \sqcap \cdots \sqcap C_n$. 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 [23]. The subsumption order $\sqsubseteq^{\emptyset}$ restricted to reduced concept descriptions is a partial order and not just a pre-order [9].

Subsumption w.r.t. the empty TBox $\emptyset$ can be recursively characterized as follows [10, Corollary 3.2]. $C \sqsubseteq^{\emptyset} D$ iff the following two conditions are satisfied:

1. For each concept name A, if $A \in \mathsf{Conj}(D)$, then $A \in \mathsf{Conj}(C)$.
2. For each existential restriction $\exists r.F$, if $\exists r.F \in \mathsf{Conj}(D)$, then there is an existential restriction $\exists r.E$ such that $\exists r.E \in \mathsf{Conj}(C)$ and $E \sqsubseteq^{\emptyset} F$.

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

Let $\mathcal{I}$ be an interpretation and X a subset of $\mathsf{Dom}(\mathcal{I})$. 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 $\mathcal{EL}_{\mathsf{si}}$ [25]. The latter DL extends $\mathcal{EL}$ with *simulation quantifiers* $\exists^{\mathsf{sim}}(\mathcal{I}, x)$ where the semantics of such concept descriptions is defined by: $y \in (\exists^{\mathsf{sim}}(\mathcal{I}, x))^{\mathcal{J}}$ if there is a simulation from $\mathcal{I}$ to $\mathcal{J}$ that contains (x, y). A *simulation* from $\mathcal{I}$ to $\mathcal{J}$ is a relation $\mathfrak{S} \subseteq \mathsf{Dom}(\mathcal{I}) \times \mathsf{Dom}(\mathcal{J})$ that satisfies the following conditions:

1. If $(x, y) \in \mathfrak{S}$ and $x \in A^{\mathcal{I}}$, then $y \in A^{\mathcal{J}}$.
2. If $(x, y) \in \mathfrak{S}$ and $(x, x') \in r^{\mathcal{I}}$, then there exists some y' such that $(x', y') \in \mathfrak{S}$ and $(y, y') \in r^{\mathcal{J}}$.

As shown in [20, Proposition 4.1.6], the MSC $X^{\mathcal{I}}$ is equivalent to $\exists^{\mathsf{sim}}(\wp(\mathcal{I}), X)$, where the *powering* $\mathcal{P}(\mathcal{I})$ has domain $\mathsf{Dom}(\wp(\mathcal{I})) := \wp(\mathsf{Dom}(\mathcal{I}))$, 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 $X \subseteq \mathsf{Dom}(\mathcal{I})$ and for each $\mathcal{EL}_{\mathsf{si}}$ 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 $\mathcal{EL}_{\mathsf{si}}$ concept description $\exists^{\mathsf{sim}}(\mathcal{I}, x)$ is always expressible in $\mathcal{EL}$ and the following recursion allows us to construct it:

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

Furthermore, the MSC $X^{\mathcal{I}}$ is equivalent to the LCS of all $\exists^{\mathsf{sim}}(\mathcal{I}, x)$ 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 $\mathcal{T} \succeq_{\mathsf{GC}} \mathcal{U}$ 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., $\mathcal{T} \succeq_{\mathsf{GC}} \mathcal{U}$ 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* $\mathcal{V} \succeq_{\mathsf{GC}} \mathcal{U}$ *but* $\mathcal{U} \not\succeq_{\mathsf{GC}} \mathcal{V}$*.*

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.

3.1 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* $C \in \mathsf{Prem}(\mathcal{T})$ *and an atom* $F \in \mathsf{Atoms}(\mathcal{P}, \mathcal{T})$ *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 [27]. Resembling the latter we define the *canonical model* $\mathcal{I}$ with domain $\mathsf{Dom}(\mathcal{I}) \coloneqq \{\, x_C \mid C \in \mathsf{Sub}(\mathcal{P}, \mathcal{T}) \,\}$ and its interpretation function is given by $A^{\mathcal{I}} \coloneqq \{\, x_C \mid C \sqsubseteq^{\mathcal{T}} A \,\}$ for each $A \in \Sigma_{\mathsf{C}}$ and $r^{\mathcal{I}} \coloneqq \{\, (x_C, x_D) \mid C \sqsubseteq^{\mathcal{T}} \exists r.D \,\}$ for each $r \in \Sigma_{\mathsf{R}}$.[4] We will now prove several technical statements involving the model $\mathcal{I}$.

Lemma I. $x_C \in C^{\mathcal{I}}$ *for each* $C \in \mathsf{Sub}(\mathcal{P}, \mathcal{T})$*.*

Proof. We show the claim by induction on the role depth of C.

– For each concept name $A \in \mathsf{Conj}(C)$, it holds that $C \sqsubseteq^{\mathcal{T}} A$ and thus $x_C \in A^{\mathcal{I}}$.

[4] In principle, this interpretation $\mathcal{I}$ is the union of the canonical models $\mathcal{I}_{C,\mathcal{T}}$ [27] for all subconcepts C occurring in $\mathcal{P}$ or $\mathcal{T}$.

– For each existential restriction $\exists r.D \in \mathsf{Conj}(C)$, we have $C \sqsubseteq^{\mathcal{T}} \exists r.D$. Since $D \in \mathsf{Sub}(\mathcal{P}, \mathcal{T})$, it follows that $(x_C, x_D) \in r^{\mathcal{I}}$. The induction hypothesis yields $x_D \in D^{\mathcal{I}}$, and we conclude that $x_C \in (\exists r.D)^{\mathcal{I}}$. □

Lemma II. *If $\mathcal{J}$ is a model of $\mathcal{T}$ and $y \in C^{\mathcal{J}}$ for some $C \in \mathsf{Sub}(\mathcal{P}, \mathcal{T})$, then there is a simulation from $\mathcal{I}$ to $\mathcal{J}$ that contains (x_C, y).*

Proof. We show that the relation $\mathfrak{S} \coloneqq \{ (x_D, z) \mid z \in D^{\mathcal{J}} \}$, which contains (x_C, y), is a simulation from $\mathcal{I}$ to $\mathcal{J}$.

– Assume $(x_D, z) \in \mathfrak{S}$ and $x_D \in A^{\mathcal{I}}$. The latter yields $D \sqsubseteq^{\mathcal{T}} A$, cf. the definition of the interpretation function, and the former yields $z \in D^{\mathcal{J}}$ by definition of $\mathfrak{S}$. Since $\mathcal{J}$ is a model of $\mathcal{T}$, we infer that $z \in A^{\mathcal{J}}$.
– Suppose that $(x_D, z) \in \mathfrak{S}$ and $(x_D, x_E) \in r^{\mathcal{I}}$, i.e., we have $z \in D^{\mathcal{J}}$ and $D \sqsubseteq^{\mathcal{T}} \exists r.E$. With $\mathcal{J}$ being a model of $\mathcal{T}$, it follows that $z \in (\exists r.E)^{\mathcal{J}}$, i.e., there is some w such that $(z, w) \in r^{\mathcal{J}}$ and $w \in E^{\mathcal{J}}$. The latter yields $(x_E, w) \in \mathfrak{S}$ and we are done. □

Lemma III. *$\mathcal{I}$ is a model of $\mathcal{T}$.*

Proof. Consider a concept inclusion $C \sqsubseteq D$ in $\mathcal{T}$, and let $x_E \in C^{\mathcal{I}}$. We first show that $E \sqsubseteq^{\mathcal{T}} D$ holds. For this purpose, let $\mathcal{J}$ be a model of $\mathcal{T}$ where $y \in E^{\mathcal{J}}$. According to Lemma II there is a simulation from $\mathcal{I}$ to $\mathcal{J}$ that contains (x_E, y). Thus $x_E \in C^{\mathcal{I}}$ implies $y \in C^{\mathcal{J}}$. Since $\mathcal{J}$ is a model of $\mathcal{T}$ and $C \sqsubseteq D$ is in $\mathcal{T}$, it follows that $y \in D^{\mathcal{J}}$.

It remains show that $x_E \in D^{\mathcal{I}}$ holds.

– For each concept name A in the top-level conjunction of D, we have $E \sqsubseteq^{\mathcal{T}} A$ and thus $x_E \in A^{\mathcal{I}}$.
– For each existential restriction $\exists r.F$ in $\mathsf{Conj}(D)$, it holds that $E \sqsubseteq^{\mathcal{T}} \exists r.F$. Since D is a subconcept of $\mathcal{T}$, also F must be in $\mathsf{Sub}(\mathcal{P}, \mathcal{T})$. It follows that $(x_E, x_F) \in r^{\mathcal{I}}$. Since $x_F \in F^{\mathcal{I}}$, we conclude that $x_E \in (\exists r.F)^{\mathcal{I}}$. □

Lemma IV. *For each subconcept $C \in \mathsf{Sub}(\mathcal{P}, \mathcal{T})$ and for each $\mathcal{EL}$ concept description E, it holds that $x_C \in E^{\mathcal{I}}$ iff $C \sqsubseteq^{\mathcal{T}} E$.*

Proof. We have already seen in the last proof that $x_C \in E^{\mathcal{I}}$ implies $C \sqsubseteq^{\mathcal{T}} E$. The converse direction follows from Lemmas I and III. □

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 $x_C \in \mathsf{Dom}(\mathcal{I})$. A* modification type *for x_C is a subset $\mathcal{K}$ of $\mathsf{Atoms}(\mathcal{P}, \mathcal{T})$ where $x_C \in K^{\mathcal{I}}$ for each $K \in \mathcal{K}$, and $K_1 \not\sqsubseteq^{\emptyset} K_2$ 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 $\{D\} \leq^{\emptyset} \mathcal{K}$ for each $D \in \mathsf{Sub}(\mathcal{P}, \mathcal{T})$ where $x_C \in D^{\mathcal{I}}$.*

Every modification type not covering a concept D w.r.t. a repair seed $\mathcal{S}$ can be enlarged to an $\mathcal{S}$-respecting modification that also does not cover D w.r.t. $\mathcal{S}$.

Lemma V. *Let $\mathcal{K}$ be a modification type for x_C such that $\{D\} \not\leq^{\mathcal{S}} \mathcal{K}$. Then there is a modification type $\mathcal{L}$ that respects $\mathcal{S}$ such that $\mathcal{K} \leq^{\emptyset} \mathcal{L}$ and $\{D\} \not\leq^{\mathcal{S}} \mathcal{L}$.*

Proof. We will construct a finite sequence of modification types $\mathcal{L}_0, \mathcal{L}_1, \ldots, \mathcal{L}_n$ for x_C each of which satisfies $\mathcal{K} \leq^{\emptyset} \mathcal{L}_i$, $\mathcal{L}_{i-1} \leq^{\emptyset} \mathcal{L}_i$ if $i > 0$, and $\{D\} \not\leq^{\mathcal{S}} \mathcal{L}_i$. To do so, we define $\mathcal{L}_0 \coloneqq \mathcal{K}$ and then construct the subsequent modification types in the following inductive manner.

- If $\mathcal{L}_i$ respects $\mathcal{S}$, then no next modification type $\mathcal{L}_{i+1}$ is constructed.
- Otherwise, there is a subconcept $E \in \mathsf{Sub}(\mathcal{P}, \mathcal{T})$ such that $x_C \in E^{\mathcal{I}}$ and $\{E\} \leq^{\mathcal{S}} \mathcal{L}_i$ but $\{E\} \not\leq^{\emptyset} \mathcal{L}_i$. So there is an atom $L \in \mathcal{L}_i$ such that $E \sqsubseteq^{\mathcal{S}} L$. Since $\{D\} \not\leq^{\mathcal{S}} \mathcal{L}_i$ we have $D \not\sqsubseteq^{\mathcal{S}} L$. It follows that $D \not\sqsubseteq^{\mathcal{S}} E$, i.e., there is an atom $L_{i+1} \in \mathsf{Conj}(E)$ where $D \not\sqsubseteq^{\mathcal{S}} L_{i+1}$. Now define the next modification type $\mathcal{L}_{i+1} \coloneqq \mathsf{Max}(\mathcal{L}_i \cup \{L_{i+1}\})$.

We show that $\mathcal{L}_{i+1}$ satisfies the two invariants.

- Consider an atom $K \in \mathcal{K}$. As $\mathcal{K} \leq^{\emptyset} \mathcal{L}_i$, there is some $L \in \mathcal{L}_i$ such that $K \sqsubseteq^{\emptyset} L$. If $L \in \mathcal{L}_{i+1}$, we are done. Otherwise, it must hold that $L \sqsubseteq^{\emptyset} L_{i+1}$, cf. the above definition of $\mathcal{L}_{i+1}$. It follows that $K \sqsubseteq^{\emptyset} L_{i+1}$. In either case K is subsumed by an atom in $\mathcal{L}_{i+1}$. We conclude that $\mathcal{K} \leq^{\emptyset} \mathcal{L}_{i+1}$.
- From $\{D\} \not\leq^{\mathcal{S}} \mathcal{L}_i$ and $D \not\sqsubseteq^{\mathcal{S}} L_{i+1}$ we infer that $\{D\} \not\leq^{\mathcal{S}} \mathcal{L}_{i+1}$.

As each $\mathcal{L}_i$ is a subset of the finite set $\mathsf{Atoms}(\mathcal{P}, \mathcal{T})$, the above rule can only be applicable finitely often. So there is a last modification type $\mathcal{L}_n$, which respects $\mathcal{S}$. Due to the invariants, $\mathcal{K} \leq^{\emptyset} \mathcal{L}_n$ and $\{D\} \not\leq^{\mathcal{S}} \mathcal{L}_n$. ☐

Note that, in the above lemma, the enlarged modification type $\mathcal{L}$ also satisfies that $\{D\} \not\leq^{\emptyset} \mathcal{L}$ since this follows if $\{D\} \not\leq^{\mathcal{S}} \mathcal{L}$.

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 $\mathsf{Dom}(\mathcal{J}_{\mathcal{S}})$ consisting of all objects $x_{C,\mathcal{K}}$ where $x_C \in \mathsf{Dom}(\mathcal{I})$ and $\mathcal{K}$ is a modification type for x_C that respects $\mathcal{S}$, and its interpretation function is defined by $A^{\mathcal{J}_{\mathcal{S}}} \coloneqq \{\, x_{C,\mathcal{K}} \mid x_C \in A^{\mathcal{I}} \text{ and } A \notin \mathcal{K} \,\}$ for each concept name $A \in \Sigma_{\mathsf{C}}$ and $r^{\mathcal{J}_{\mathcal{S}}} \coloneqq \{\, (x_{C,\mathcal{K}}, x_{D,\mathcal{L}}) \mid (x_C, x_D) \in r^{\mathcal{I}} \text{ and } \mathsf{Succ}(\mathcal{K}, r, x_D) \leq^{\emptyset} \mathcal{L} \,\}$ for each role name $r \in \Sigma_{\mathsf{R}}$, where $\mathsf{Succ}(\mathcal{K}, r, x_D) \coloneqq \{\, E \mid \exists r.E \in \mathcal{K} \text{ and } x_D \in E^{\mathcal{I}} \,\}$.*

The next lemma shows structural properties of the induced countermodel $\mathcal{J}_{\mathcal{S}}$.

Lemma VI. *Consider an object $x_{C,\mathcal{K}} \in \mathsf{Dom}(\mathcal{J}_{\mathcal{S}})$ and an $\mathcal{EL}$ concept description E.*

1. *If $x_{C,\mathcal{K}} \in E^{\mathcal{J}_{\mathcal{S}}}$, then $x_C \in E^{\mathcal{I}}$ and $\{E\} \not\leq^{\emptyset} \mathcal{K}$.*
2. *If $x_C \in E^{\mathcal{I}}$ and $\{E\} \not\leq^{\mathcal{S}} \mathcal{K}$, then $x_{C,\mathcal{K}} \in E^{\mathcal{J}_{\mathcal{S}}}$.*

Proof. We start with proving the first statement. It is easy to verify that the relation $\{\ (x_{C,\mathcal{K}}, x_C) \mid x_{C,\mathcal{K}} \in \mathsf{Dom}(\mathcal{J}_\mathcal{S})\ \}$ is a simulation from the induced countermodel $\mathcal{J}_\mathcal{S}$ to the canonical model $\mathcal{I}$. Thus $x_{C,\mathcal{K}} \in E^{\mathcal{J}_\mathcal{S}}$ implies $x_C \in E^{\mathcal{I}}$.

Next, we show by contraposition that $x_{C,\mathcal{K}} \in E^{\mathcal{J}_\mathcal{S}}$ implies $\{E\} \not\leq^\emptyset \mathcal{K}$. From the precondition $\{E\} \leq^\emptyset \mathcal{K}$ it follows that there is some atom $K \in \mathcal{K}$ such that $E \sqsubseteq^\emptyset K$. We proceed with an induction on K.

- The case where K is a concept name A is obvious: then $x_{C,\mathcal{K}} \notin A^{\mathcal{J}_\mathcal{S}}$ by Definition 8, and so $E \sqsubseteq^\emptyset A$ implies $x_{C,\mathcal{K}} \notin E^{\mathcal{J}_\mathcal{S}}$.
- Now assume that K is an existential restriction $\exists r.F$. Further let $x_{D,\mathcal{L}}$ be an r-successor of $x_{C,\mathcal{K}}$, i.e., $\mathsf{Succ}(\mathcal{K}, r, x_D) \leq^\emptyset \mathcal{L}$. If $x_D \notin F^{\mathcal{I}}$, then the existence of the above simulation implies $x_{D,\mathcal{K}} \notin F^{\mathcal{J}_\mathcal{S}}$. Otherwise we have $\{F\} \leq^\emptyset \mathcal{L}$ and so there is an atom $L \in \mathcal{L}$ such that $F \sqsubseteq^\emptyset L$. Since the role depth of L is smaller than the role depth of K, it follows by induction hypothesis that $x_{D,\mathcal{L}} \notin F^{\mathcal{J}_\mathcal{S}}$. We conclude that $x_{C,\mathcal{K}} \notin (\exists r.F)^{\mathcal{J}_\mathcal{S}}$, and so $E \sqsubseteq^\emptyset \exists r.F$ yields $x_{C,\mathcal{K}} \notin E^{\mathcal{J}_\mathcal{S}}$.

Last, we show the second statement. Therefore assume that $x_C \in E^{\mathcal{I}}$ and $\{E\} \not\leq^\mathcal{S} \mathcal{K}$. We continue with an induction on E.

- The case where E is the top concept $\top$ is trivial.
- Assume that E is a concept name A. Since $\{A\} \not\leq^\mathcal{S} \mathcal{K}$ implies $\{A\} \not\leq^\emptyset \mathcal{K}$, it follows that $A \notin \mathcal{K}$ and thus $x_{C,\mathcal{K}} \in A^{\mathcal{J}_\mathcal{S}}$ by Definition 6.
- If E is a conjunction $E_1 \sqcap E_2$, then $\{E_1 \sqcap E_2\} \not\leq^\mathcal{S} \mathcal{K}$ implies $\{E_1\} \not\leq^\mathcal{S} \mathcal{K}$ and $\{E_2\} \not\leq^\mathcal{S} \mathcal{K}$. The induction hypothesis yields $x_{C,\mathcal{K}} \in E_1^{\mathcal{J}_\mathcal{S}}$ and $x_{C,\mathcal{K}} \in E_2^{\mathcal{J}_\mathcal{S}}$, and thus $x_{C,\mathcal{K}} \in (E_1 \sqcap E_2)^{\mathcal{J}_\mathcal{S}}$.
- Now consider the case where E is an existential restriction $\exists r.F$. Since $x_C \in (\exists r.F)^{\mathcal{I}}$, there is some x_D such that $(x_C, x_D) \in r^{\mathcal{I}}$ and $x_D \in F^{\mathcal{I}}$. We must find a modification type $\mathcal{L}$ for x_D that respects $\mathcal{S}$ and such that $\mathsf{Succ}(\mathcal{K}, r, x_D) \leq^\emptyset \mathcal{L}$ and $\{F\} \not\leq^\mathcal{S} \mathcal{L}$. Then $(x_{C,\mathcal{K}}, x_{D,\mathcal{L}}) \in r^{\mathcal{J}_\mathcal{S}}$ and the induction hypothesis yields $x_{D,\mathcal{L}} \in F^{\mathcal{J}_\mathcal{S}}$, which together implies that $x_{C,\mathcal{K}} \in (\exists r.F)^{\mathcal{J}_\mathcal{S}}$.
  Since $\{\exists r.F\} \not\leq^\mathcal{S} \mathcal{K}$, we specifically have $\{F\} \not\leq^\mathcal{S} \mathsf{Succ}(\mathcal{K}, r, x_D)$. This means that, for each $\exists r.G \in \mathcal{K}$ where $x_D \in G^{\mathcal{I}}$, there is a top-level conjunct $L_{\exists r.G} \in \mathsf{Conj}(G)$ such that $F \not\sqsubseteq^\mathcal{S} L_{\exists r.G}$. Now consider the following modification type for x_D:

  $$\mathcal{L}_0 \coloneqq \mathsf{Max}\{\ L_{\exists r.G} \mid \exists r.G \in \mathcal{K} \text{ and } x_D \in G^{\mathcal{I}}\ \}.$$

  It holds that $\mathsf{Succ}(\mathcal{K}, r, x_D) \leq^\emptyset \mathcal{L}_0$ and $\{F\} \not\leq^\mathcal{S} \mathcal{L}_0$. According to Lemma V there is a modification type $\mathcal{L}$ for x_D that respects $\mathcal{S}$ and satisfies $\mathcal{L}_0 \leq^\emptyset \mathcal{L}$ as well as $\{F\} \not\leq^\mathcal{S} \mathcal{L}$. The former implies that $\mathsf{Succ}(\mathcal{K}, r, x_D) \leq^\emptyset \mathcal{L}$ as needed. □

Since $\mathcal{S}$ complies with $\mathcal{P}$, we can show by means of the last two lemmas that $\mathcal{J}_\mathcal{S}$ is indeed a countermodel to $\mathcal{P}$.

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

Proof. Let $C \sqsubseteq D$ be a concept inclusion in the repair request $\mathcal{P}$. It follows that the canonical model $\mathcal{I}$ contains the object x_C. Since $\mathcal{S}$ does not entail $C \sqsubseteq D$, there is a top-level conjunct $E \in \mathsf{Conj}(D)$ such that $C \not\sqsubseteq^{\mathcal{S}} E$. Since $\mathcal{T}$ entails $C \sqsubseteq D$, it follows by Lemma IV that $x_C \in E^{\mathcal{I}}$ and so $\{E\}$ is a modification type for x_C. By Lemma V there is a modification type $\mathcal{K}$ for x_C that respects $\mathcal{S}$ and fulfills $\{E\} \leq^{\emptyset} \mathcal{K}$ and $\{C\} \not\leq^{\mathcal{S}} \mathcal{K}$. Of course, the former implies $\{D\} \leq^{\emptyset} \mathcal{K}$. Now the object $x_{C,\mathcal{K}}$ is in the domain of the induced countermodel $\mathcal{J}_{\mathcal{S}}$. With Lemma VI we infer that $x_{C,\mathcal{K}} \in C^{\mathcal{J}_{\mathcal{S}}}$ and $x_{C,\mathcal{K}} \notin D^{\mathcal{J}_{\mathcal{S}}}$, i.e., $\mathcal{J}_{\mathcal{S}}$ does not satisfy $C \sqsubseteq D$. □

3.2 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*

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

The following lemma shows that $\mathsf{rep}_{\mathsf{GC}}(\mathcal{T}, \mathcal{S})$ 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. $\mathsf{rep}_{\mathsf{GC}}(\mathcal{T}, \mathcal{S}) \succeq_{\mathsf{GC}} \mathcal{U}$ *iff* $\mathcal{T} \succeq_{\mathsf{GC}} \mathcal{U}$ *and* $\mathcal{J}_{\mathcal{S}} \models \mathcal{U}$

Proof. We first prove the only-if direction, so let $\mathsf{rep}_{\mathsf{GC}}(\mathcal{T}, \mathcal{S}) \succeq_{\mathsf{GC}} \mathcal{U}$. Consider a concept inclusion $C \sqsubseteq D$ in $\mathcal{U}$. So there is a concept inclusion $E \sqsubseteq F$ in $\mathsf{rep}_{\mathsf{GC}}(\mathcal{T}, \mathcal{S})$ such that $C = E$ and $F \sqsubseteq^{\emptyset} D$. By Definition 8, F must be of the form $F' \vee E^{\mathcal{J}_{\mathcal{S}}\mathcal{J}_{\mathcal{S}}}$ where $\mathcal{T}$ contains $E \sqsubseteq F'$. It follows that $F' \sqsubseteq^{\emptyset} F$ and thus $F' \sqsubseteq^{\emptyset} D$. We further conclude that $E^{\mathcal{J}_{\mathcal{S}}\mathcal{J}_{\mathcal{S}}} \sqsubseteq^{\emptyset} F$, which means that $\mathcal{J}_{\mathcal{S}} \models E \sqsubseteq F$. From $C = E$ and $F \sqsubseteq^{\emptyset} D$ we thus infer that $\mathcal{J}_{\mathcal{S}} \models C \sqsubseteq D$.

Next, we show the if direction. Assume $\mathcal{T} \succeq_{\mathsf{GC}} \mathcal{U}$ and $\mathcal{J}_{\mathcal{S}} \models \mathcal{U}$, and consider a concept inclusion $C \sqsubseteq D$ in $\mathcal{U}$. Since $\mathcal{T} \succeq_{\mathsf{GC}} \mathcal{U}$, there is a concept inclusion $E \sqsubseteq F$ in $\mathcal{T}$ where $C = E$ and $F \sqsubseteq^{\emptyset} D$. From $\mathcal{J}_{\mathcal{S}} \models \mathcal{U}$ we obtain that $\mathcal{J}_{\mathcal{S}} \models C \sqsubseteq D$ and thus $E^{\mathcal{J}_{\mathcal{S}}\mathcal{J}_{\mathcal{S}}} \equiv^{\emptyset} C^{\mathcal{J}_{\mathcal{S}}\mathcal{J}_{\mathcal{S}}} \sqsubseteq^{\emptyset} D$. It follows that $F \vee E^{\mathcal{J}_{\mathcal{S}}\mathcal{J}_{\mathcal{S}}} \sqsubseteq^{\emptyset} D$. According to Definition 8, $\mathsf{rep}_{\mathsf{GC}}(\mathcal{T}, \mathcal{S})$ contains the CI $E \sqsubseteq F \vee E^{\mathcal{J}_{\mathcal{S}}\mathcal{J}_{\mathcal{S}}}$. □

As $\mathsf{rep}_{\mathsf{GC}}(\mathcal{T}, \mathcal{S})$ is a GC-weakening of itself, we infer that $\mathcal{J}_{\mathcal{S}}$ is a model of $\mathsf{rep}_{\mathsf{GC}}(\mathcal{T}, \mathcal{S})$. According to Proposition 7, $\mathcal{J}_{\mathcal{S}}$ is a countermodel to $\mathcal{P}$, and so $\mathsf{rep}_{\mathsf{GC}}(\mathcal{T}, \mathcal{S})$ complies with $\mathcal{P}$. It is further easy to see that $\mathsf{rep}_{\mathsf{GC}}(\mathcal{T}, \mathcal{S})$ 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* $\mathsf{rep}_{\mathsf{GC}}(\mathcal{T}, \mathcal{S})$ *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., $\mathsf{rep}_{\mathsf{GC}}(\mathcal{T}, \emptyset)$ 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 $\mathsf{rep}_{\mathsf{GC}}(\mathcal{T}, \mathcal{S})$ 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* $\mathsf{rep}_{\mathsf{GC}}(\mathcal{T},\mathcal{S}) \succeq_{\mathsf{GC}} \mathcal{U}$.

Proof. Suppose that $\mathcal{U}$ is a GC-repair of $\mathcal{T}$ for $\mathcal{P}$. Let $\mathcal{S}_{\mathcal{U}}^0 \coloneqq \emptyset$ and

$$\mathcal{S}_{\mathcal{U}}^{n+1} \coloneqq \{\, C \sqsubseteq F \mid C \sqsubseteq D' \in \mathcal{U},\ F \in \mathsf{Atoms}(\mathcal{P},\mathcal{T}),\ \text{and } D' \sqsubseteq^{\mathcal{S}_{\mathcal{U}}^n} F \,\}$$

for each number $n \geq 0$. It is easy to see that $\mathcal{S}_{\mathcal{U}}^n \subseteq \mathcal{S}_{\mathcal{U}}^{n+1}$ always holds. Since the TBox $\mathcal{U}$ as well as the set $\mathsf{Atoms}(\mathcal{P},\mathcal{T})$ are finite, there must be an index n such that the subset inclusion is actually an equality — then define $\mathcal{S}_{\mathcal{U}}^* \coloneqq \mathcal{S}_{\mathcal{U}}^n$. This fixed point satisfies the equation $\mathcal{S}_{\mathcal{U}}^* = \{\, C \sqsubseteq F \mid C \sqsubseteq D' \in \mathcal{U},\ F \in \mathsf{Atoms}(\mathcal{P},\mathcal{T})$, and $D' \sqsubseteq^{\mathcal{S}_{\mathcal{U}}^*} F \,\}$. It has the important property that $D' \sqsubseteq^{\mathcal{S}_{\mathcal{U}}^*} F$ implies $C \sqsubseteq^{\mathcal{S}_{\mathcal{U}}^*} F$ for each concept inclusion $C \sqsubseteq D' \in \mathcal{U}$ and each atom $F \in \mathsf{Atoms}(\mathcal{P},\mathcal{T})$. Thus, $\{D'\} \leq^{\mathcal{S}_{\mathcal{U}}^*} \mathcal{K}$ implies $\{C\} \leq^{\mathcal{S}_{\mathcal{U}}^*} \mathcal{K}$ for each modification type $\mathcal{K}$ — we will use this property later.

We show that $\mathcal{U}$ entails $\mathcal{S}_{\mathcal{U}}^*$, namely by induction. $\mathcal{U}$ trivially entails $\mathcal{S}_{\mathcal{U}}^0$. Now assume that $\mathcal{U}$ entails $\mathcal{S}_{\mathcal{U}}^n$, and consider a concept inclusion $C \sqsubseteq F$ in $\mathcal{S}_{\mathcal{U}}^{n+1}$, i.e., $C \sqsubseteq D' \in \mathcal{U}$, $F \in \mathsf{Atoms}(\mathcal{P},\mathcal{T})$, and $D' \sqsubseteq^{\mathcal{S}_{\mathcal{U}}^n} F$. From the latter we infer that $D' \sqsubseteq^{\mathcal{U}} F$, which together with $C \sqsubseteq D' \in \mathcal{U}$ implies that $C \sqsubseteq^{\mathcal{U}} F$ as needed.

Since $\mathcal{U}$ does not entail any concept inclusion in the repair request $\mathcal{P}$, also $\mathcal{S}_{\mathcal{U}}^*$ complies with $\mathcal{P}$. From the precondition $\mathcal{T} \succeq_{\mathsf{GC}} \mathcal{U}$ it follows that $\mathcal{T} \models \mathcal{U}$ and thus $\mathcal{T} \models \mathcal{S}_{\mathcal{U}}^*$. We conclude that $\mathcal{S}_{\mathcal{U}}^*$ is a repair seed.

Next, we show that $\mathsf{rep}_{\mathsf{GC}}(\mathcal{T},\mathcal{S}_{\mathcal{U}}^*) \succeq_{\mathsf{GC}} \mathcal{U}$. Since $\mathcal{U}$ is a GC-weakening of $\mathcal{T}$, it suffices to show that the induced countermodel $\mathcal{J}_{\mathcal{S}_{\mathcal{U}}^*}$ is a model of $\mathcal{U}$, cf. Lemma 9.

Consider a concept inclusion $C \sqsubseteq D'$ in $\mathcal{U}$. Then there is a concept inclusion $C \sqsubseteq D$ in $\mathcal{T}$ such that $D \sqsubseteq^{\emptyset} D'$. Further assume $x_{E,\mathcal{K}} \in C^{\mathcal{J}_{\mathcal{S}_{\mathcal{U}}^*}}$. Lemma VI implies $E \sqsubseteq^{\mathcal{T}} C$ and $\{C\} \not\leq^{\emptyset} \mathcal{K}$. We must show that $x_{E,\mathcal{K}} \in (D')^{\mathcal{J}_{\mathcal{S}_{\mathcal{U}}^*}}$ holds as well, which according to Lemma VI is implied by $E \sqsubseteq^{\mathcal{T}} D'$ and $\{D'\} \not\leq^{\mathcal{S}_{\mathcal{U}}^*} \mathcal{K}$. The former follows from $E \sqsubseteq^{\mathcal{T}} C$, $C \sqsubseteq D \in \mathcal{T}$, and $D \sqsubseteq^{\emptyset} D'$. Since $C \in \mathsf{Sub}(\mathcal{P},\mathcal{T})$ and $\mathcal{K}$ respects $\mathcal{S}_{\mathcal{U}}^*$, we conclude from $\{C\} \not\leq^{\emptyset} \mathcal{K}$ that $\{C\} \not\leq^{\mathcal{S}_{\mathcal{U}}^*} \mathcal{K}$. As already shown above, it follows that $\{D'\} \not\leq^{\mathcal{S}_{\mathcal{U}}^*} \mathcal{K}$. ◻

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. $\succeq_{\mathsf{GC}}$. 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 $\{\exists r.A \sqsubseteq \exists r.(P_1 \sqcap Q_1 \sqcap \cdots \sqcap P_n \sqcap Q_n),\ P_1 \sqcap Q_1 \sqsubseteq B,\ \ldots,\ P_n \sqcap Q_n \sqsubseteq B\}$ has the optimal GC-repair $\{\exists r.A \sqsubseteq \bigsqcap\{\, \exists r.(X_1 \sqcap \cdots \sqcap X_n) \mid X_i \in \{P_i, Q_i\} \text{ for each } i \in \{1,\ldots,n\} \,\}$, $P_1 \sqcap Q_1 \sqsubseteq B,\ \ldots,\ P_n \sqcap Q_n \sqsubseteq B\}$. It has exponential size.

3.3 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 $D \vee C^{\mathcal{J}_\mathcal{S}\mathcal{J}_\mathcal{S}}$ can be obtained as the product of the $\mathcal{EL}$ concept description D and the $\mathcal{EL}_\mathsf{si}$ concept description $\exists^\mathsf{sim}(\wp(\mathcal{J}_\mathcal{S}), C^{\mathcal{J}_\mathcal{S}})$, or alternatively as the product of D and all $\exists^\mathsf{sim}(\mathcal{J}_\mathcal{S}, x_{E,\mathcal{K}})$ where $x_{E,\mathcal{K}} \in C^{\mathcal{J}_\mathcal{S}}$. However, computing the induced GC-repair $\mathsf{rep}_\mathsf{GC}(\mathcal{T}, \mathcal{S})$ in this way is very inefficient since $\mathcal{J}_\mathcal{S}$ has exponential size.

The first important observation is that the concept description $C^{\mathcal{J}_\mathcal{S}\mathcal{J}_\mathcal{S}}$ is already equivalent to $\exists^\mathsf{sim}(\mathcal{J}_\mathcal{S}, x_{C,\mathcal{S}[C]})$ where $\mathcal{S}[C]$ is the largest modification type for x_C that respects $\mathcal{S}$ and does not contain an atom subsuming C. This follows from the fact that there is a simulation on $\mathcal{J}_\mathcal{S}$ that contains the pair $(x_{C,\mathcal{S}[C]}, x_{E,\mathcal{K}})$ for each object $x_{E,\mathcal{K}}$ in the extension $C^{\mathcal{J}_\mathcal{S}}$. We verify this observation in Lemmas VII–IX.

Lemma VII. $C^{\mathcal{J}_\mathcal{S}\mathcal{J}_\mathcal{S}} \equiv^\emptyset \bigvee\{\, \exists^\mathsf{sim}(\mathcal{J}_\mathcal{S}, x_{E,\mathcal{K}}) \mid E \sqsubseteq^\mathcal{T} C$ *and* $\{C\} \not\leq^\mathcal{S} \mathcal{K} \,\}$

Proof. Since $C \in \mathsf{Sub}(\mathcal{P}, \mathcal{T})$, Lemma VI yields $C^{\mathcal{J}_\mathcal{S}} = \{\, x_{E,\mathcal{K}} \mid E \sqsubseteq^\mathcal{T} C$ and $\{C\} \not\leq^\mathcal{S} \mathcal{K} \,\}$. As already mentioned in Section 2, the MSC $X^{\mathcal{J}_\mathcal{S}}$ is equivalent to the LCS of all $\exists^\mathsf{sim}(\mathcal{J}_\mathcal{S}, x)$ where $x \in X$. The claim follows for $X \coloneqq C^{\mathcal{J}_\mathcal{S}}$. □

Recall from Lemma IV that, for each subconcept $C \in \mathsf{Sub}(\mathcal{P}, \mathcal{T})$ and for each $\mathcal{EL}$ concept description E, it holds that $x_C \in E^\mathcal{I}$ iff $C \sqsubseteq^\mathcal{T} E$. We will implicitly use this equivalence in the following. We specifically have $\mathsf{Succ}(\mathcal{K}, r, x_D) = \{\, E \mid \exists r.E \in \mathcal{K}$ and $D \sqsubseteq^\mathcal{T} E \,\}$, and we will sometimes write $\mathsf{Succ}(\mathcal{K}, r, D)$ for this set.

Lemma VIII. *The relation* $\mathfrak{S} \coloneqq \{\, (x_{E,\mathcal{K}}, x_{F,\mathcal{L}}) \mid F \sqsubseteq^\mathcal{T} E$ *and* $\mathcal{L}\restriction_E \leq^\emptyset \mathcal{K} \,\}$ *is a simulation on* $\mathcal{J}_\mathcal{S}$*, where* $\mathcal{L}\restriction_E \coloneqq \{\, L \mid L \in \mathcal{L}$ *and* $E \sqsubseteq^\mathcal{T} L \,\}$.

Proof. Let $x_{E,\mathcal{K}} \in A^{\mathcal{J}_\mathcal{S}}$, i.e., $E \sqsubseteq^\mathcal{T} A$ and $A \notin \mathcal{K}$. It follows that $F \sqsubseteq^\mathcal{T} A$ and $A \notin \mathcal{L}\restriction_E$. We further obtain that $A \notin \mathcal{L}$ and thus $x_{F,\mathcal{L}} \in A^{\mathcal{J}_\mathcal{S}}$.

Assume $(x_{E,\mathcal{K}}, x_{G,\mathcal{M}}) \in r^{\mathcal{J}_\mathcal{S}}$, i.e., $E \sqsubseteq^\mathcal{T} \exists r.G$ and $\mathsf{Succ}(\mathcal{K}, r, G) \leq^\emptyset \mathcal{M}$. Then we infer that also $F \sqsubseteq^\mathcal{T} \exists r.G$ and $\mathsf{Succ}(\mathcal{L}\restriction_E, r, G) \leq^\emptyset \mathcal{M}$.

Let $\exists r.H \in \mathcal{L}$ where $E \not\sqsubseteq^\mathcal{T} \exists r.H$. Then $G \not\sqsubseteq^\mathcal{T} H$, and thus $H \notin \mathsf{Succ}(\mathcal{L}, r, G)$. We thus obtain that $\mathsf{Succ}(\mathcal{L}\restriction_E, r, G) = \mathsf{Succ}(\mathcal{L}, r, G)$, which yields that $\mathsf{Succ}(\mathcal{L}, r, G) \leq^\emptyset \mathcal{M}$. We conclude that $(x_{F,\mathcal{L}}, x_{G,\mathcal{M}}) \in r^{\mathcal{J}_\mathcal{S}}$. Clearly, $\mathfrak{S}$ contains the pair $(x_{G,\mathcal{M}}, x_{G,\mathcal{M}})$. □

Lemma IX. $C^{\mathcal{J}_\mathcal{S}\mathcal{J}_\mathcal{S}} \equiv^\emptyset \exists^\mathsf{sim}(\mathcal{J}_\mathcal{S}, x_{C,\mathcal{S}[C]})$ *where*

$$\mathcal{S}[C] \coloneqq \mathsf{Max}\{\, K \mid K \in \mathsf{Atoms}(\mathcal{P}, \mathcal{T}),\ C \not\sqsubseteq^\mathcal{S} K,\ \textit{and } C \sqsubseteq^\mathcal{T} K \,\}.$$

Proof. $\mathcal{S}[C]$ clearly is a modification type for x_C. In order to verify that $\mathcal{S}[C]$ respects $\mathcal{S}$, consider a subconcept $F \in \mathsf{Sub}(\mathcal{P}, \mathcal{T})$ such that $C \sqsubseteq^\mathcal{T} F$ and $\{F\} \leq^\mathcal{S} \mathcal{S}[C]$. So there is an atom K in $\mathcal{S}[C]$ where $F \sqsubseteq^\mathcal{S} K$. It follows that $C \not\sqsubseteq^\mathcal{S} F$, and so either F itself is in $\mathcal{S}[C]$ or F is subsumed by an atom in $\mathcal{S}[C]$, i.e., $\{F\} \leq^\emptyset \mathcal{S}[C]$.

Next, we show that $\exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{E,\mathcal{K}}) \sqsubseteq^{\emptyset} \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{C,\mathcal{S}[C]})$ holds for each $x_{E,\mathcal{K}} \in C^{\mathcal{J}_{\mathcal{S}}}$, namely by proving that the simulation $\mathfrak{S}$ in Lemma VIII contains the pair $(x_{C,\mathcal{S}[C]}, x_{E,\mathcal{K}})$ [20, Proposition 3.4.2]. So consider an object $x_{E,\mathcal{K}} \in C^{\mathcal{J}_{\mathcal{S}}}$. We already know from the proof of Lemma VII that $E \sqsubseteq^{\mathcal{T}} C$ and $\{C\} \not\leq^{\mathcal{S}} \mathcal{K}$. It remains to verify that $\mathcal{K}\restriction_C \leq^{\emptyset} \mathcal{S}[C]$. Suppose that K is an atom in $\mathcal{K}$ where $C \sqsubseteq^{\mathcal{T}} K$. Since $\{C\} \not\leq^{\mathcal{S}} \mathcal{K}$, we further have $C \not\sqsubseteq^{\mathcal{S}} K$. We conclude that $\mathcal{S}[C]$ must contain either K itself or an atom subsuming K. ☐

Secondly, in order to compute the LCS $D \vee \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{C,\mathcal{S}[C]})$ 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. We show this in Lemmas XI–XIV.

Definition X. *We say that $\mathcal{L}$ is* tolerated *by $x_{C,\mathcal{K}}$ if each atom L in $\mathcal{L}$ satisfies $C \not\sqsubseteq^{\mathcal{T}} L$ or $\{L\} \leq^{\emptyset} \mathcal{K}$.*

Lemma XI. *If $\mathcal{M}$ is covered by $\mathcal{L}$ and $\mathcal{L}$ is tolerated by $x_{C,\mathcal{K}}$, then $\mathcal{M}$ is tolerated by $x_{C,\mathcal{K}}$.*

Proof. Consider an atom $M \in \mathcal{M}$. Since $\mathcal{M} \leq^{\emptyset} \mathcal{L}$, there is an atom $L \in \mathcal{L}$ such that $M \sqsubseteq^{\emptyset} L$. As $\mathcal{L}$ is tolerated by $x_{C,\mathcal{K}}$, we have $C \not\sqsubseteq^{\mathcal{T}} L$ or $\{L\} \leq^{\emptyset} \mathcal{K}$. In the latter case we immediately infer that $\{M\} \leq^{\emptyset} \mathcal{K}$. Otherwise, $C \sqsubseteq^{\mathcal{T}} M$ cannot hold as it would yield, using $M \sqsubseteq^{\emptyset} L$, a contradiction to $C \not\sqsubseteq^{\mathcal{T}} L$. ☐

Lemma XII. *If $\mathcal{N}$ is a modification type for x_F that is tolerated by $x_{E,\mathcal{M}}$, then there is a modification type $\mathcal{N}'$ for x_F that respects $\mathcal{S}$, is tolerated by $x_{E,\mathcal{M}}$, and satisfies $\mathcal{N} \leq^{\emptyset} \mathcal{N}'$.*

Proof. We initialize $\mathcal{N}' \coloneqq \mathcal{N}$ and then exhaustively apply the following rule. The invariant is that $\mathcal{N}'$ is always tolerated by $x_{E,\mathcal{M}}$.

Extension Rule. If there is a subconcept $G \in \mathsf{Sub}(\mathcal{P}, \mathcal{T})$ where $F \sqsubseteq^{\mathcal{T}} G$ and $\{G\} \leq^{\mathcal{S}} \mathcal{N}'$ but $\{G\} \not\leq^{\emptyset} \mathcal{N}'$, then extend $\mathcal{N}'$ as follows; otherwise this rule is not applicable. We first infer that there is an atom $N \in \mathcal{N}'$ such that $G \sqsubseteq^{\mathcal{S}} N$.

- If $E \sqsubseteq^{\mathcal{T}} G$, then N must be subsumed by an atom in $\mathcal{M}$. (This follows from the invariant: if $\{N\}$ was not covered by $\mathcal{M}$, then it would hold that $E \not\sqsubseteq^{\mathcal{T}} N$, but from $E \sqsubseteq^{\mathcal{T}} G \sqsubseteq^{\mathcal{S}} N$ we could infer the contradiction that $E \sqsubseteq^{\mathcal{T}} N$.) This means that $\{G\} \leq^{\mathcal{S}} \mathcal{M}$. Since $\mathcal{M}$ respects $\mathcal{S}$, there is an atom $M \in \mathcal{M}$ such that $G \sqsubseteq^{\emptyset} M$. Replace $\mathcal{N}'$ with $\mathsf{Max}(\mathcal{N}' \cup \{M\})$.
- If $E \not\sqsubseteq^{\mathcal{T}} G$, then there is an atom $G' \in \mathsf{Conj}(G)$ such that $E \not\sqsubseteq^{\mathcal{T}} G'$. Replace $\mathcal{N}'$ with $\mathsf{Max}(\mathcal{N}' \cup \{G'\})$.

Since there are only finitely many atoms in $\mathsf{Atoms}(\mathcal{P}, \mathcal{T})$, rule application must terminate after finitely many steps. It is easy to verify that the final set $\mathcal{N}'$ is a modification type for x_F, respects $\mathcal{S}$ (since the Extension Rule is not applicable), is tolerated by $x_{E,\mathcal{M}}$ (due to the invariant), and covers $\mathcal{N}$. ☐

Lemma XIII. *If $x_{C,\mathcal{K}}$ and $x_{D,\mathcal{L}}$ are elements of $\mathsf{Dom}(\mathcal{J}_{\mathcal{S}})$ such that $\mathcal{L}$ is tolerated by $x_{C,\mathcal{K}}$, then $D \vee \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{D,\mathcal{L}})$ is subsumed by $D \vee \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{C,\mathcal{K}})$.*

Proof. We show the claim by induction on D.

Consider a concept name A in the top-level conjunction of $D \vee \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{C,\mathcal{K}})$, i.e., $A \in \mathsf{Conj}(D)$, and $C \sqsubseteq^{\mathcal{T}} A$, and $A \notin \mathcal{K}$. Since $\mathcal{L}$ is tolerated by $x_{C,\mathcal{K}}$, it follows that $A \notin \mathcal{L}$. Clearly, $A \in \mathsf{Conj}(D)$ implies $D \sqsubseteq^{\mathcal{T}} A$. We conclude that $x_{D,\mathcal{L}} \in A^{\mathcal{J}_{\mathcal{S}}}$, i.e., A is a top-level conjunct of $D \vee \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{D,\mathcal{L}})$.

Next, let $\exists r.(F \vee \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{E,\mathcal{M}}))$ be a top-level conjunct of $D \vee \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{C,\mathcal{K}})$, i.e., $\exists r.F \in \mathsf{Conj}(D)$, $C \sqsubseteq^{\mathcal{T}} \exists r.E$, and $\mathsf{Succ}(\mathcal{K}, r, E) \leq^{\emptyset} \mathcal{M}$.

Since $\exists r.F$ is a top-level conjunct of D, we also have $D \sqsubseteq^{\mathcal{T}} \exists r.F$. We will construct a modification type $\mathcal{N}$ for x_F that is tolerated by $x_{E,\mathcal{M}}$ and such that $(x_{D,\mathcal{L}}, x_{F,\mathcal{N}}) \in r^{\mathcal{J}_{\mathcal{S}}}$. Then the induction hypothesis yields that $F \vee \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{E,\mathcal{M}})$ subsumes $F \vee \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{F,\mathcal{N}})$. It then further follows that $D \vee \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{D,\mathcal{L}})$ contains the top-level conjunct $\exists r.(F \vee \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{F,\mathcal{N}}))$ and the latter subsumes $\exists r.(F \vee \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{E,\mathcal{M}}))$.

Initialize $\mathcal{N} \coloneqq \emptyset$. We first add enough atoms to $\mathcal{N}$ such that it covers $\mathsf{Succ}(\mathcal{L}, r, F)$. Suppose that $H \in \mathsf{Succ}(\mathcal{L}, r, F)$, i.e., $\exists r.H \in \mathcal{L}$ where $F \sqsubseteq^{\mathcal{T}} H$. Since $\mathcal{L}$ is tolerated by $x_{C,\mathcal{K}}$, we have $C \not\sqsubseteq^{\mathcal{T}} \exists r.H$ or $\{\exists r.H\} \leq^{\emptyset} \mathcal{K}$. We proceed with a case distinction.

- Assume that $E \sqsubseteq^{\mathcal{T}} H$. Since $C \sqsubseteq^{\mathcal{T}} \exists r.E$, it follows that $C \sqsubseteq^{\mathcal{T}} \exists r.H$, which further implies $\{\exists r.H\} \leq^{\emptyset} \mathcal{K}$, i.e., there is an atom $\exists r.H' \in \mathcal{K}$ such that $H \sqsubseteq^{\emptyset} H'$. We infer that $H' \in \mathsf{Succ}(\mathcal{K}, r, E)$ and, since the latter is covered by $\mathcal{M}$, there is an atom $M \in \mathcal{M}$ such that $H' \sqsubseteq^{\emptyset} M$, i.e., it also holds that $H \sqsubseteq^{\emptyset} M$. Replace $\mathcal{N}$ with $\mathsf{Max}(\mathcal{N} \cup \{M\})$.
- In the remaining case we have $E \not\sqsubseteq^{\mathcal{T}} H$. So there is an atom $H' \in \mathsf{Conj}(H)$ such that $E \not\sqsubseteq^{\mathcal{T}} H'$. Replace $\mathcal{N}$ with $\mathsf{Max}(\mathcal{N} \cup \{H'\})$.

The resulting set $\mathcal{N}$ is a modification type for x_F, is tolerated by $x_{E,\mathcal{M}}$, but it need not respect $\mathcal{S}$. According to Lemma XII, we can extend $\mathcal{N}$ to an $\mathcal{S}$-respecting modification type and we are done. □

Lemma XIV. $D \vee \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{C,\mathcal{S}[C]}) \equiv^{\emptyset} D \vee \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{D,\mathcal{S}[C \sqsubseteq D]})$ *where*

$$\mathcal{S}[C \sqsubseteq D] \coloneqq \mathcal{S}[C]\restriction_D = \mathsf{Max}\{\, K \mid K \in \mathsf{Atoms}(\mathcal{P}, \mathcal{T}),\ C \not\sqsubseteq^{\mathcal{S}} K,\ \textit{and}\ D \sqsubseteq^{\mathcal{T}} K \,\}.$$

Proof. It is easy to see that $\mathcal{S}[C]\restriction_D$ is a modification type for x_D. It remains to show that it respects $\mathcal{S}$. Suppose a subconcept $E \in \mathsf{Sub}(\mathcal{P}, \mathcal{T})$ such that $D \sqsubseteq^{\mathcal{T}} E$ and $\{E\} \leq^{\mathcal{S}} \mathcal{S}[C]\restriction_D$. Since $\mathcal{S}[C]$ respects $\mathcal{S}$, it follows that $\{E\} \leq^{\emptyset} \mathcal{S}[C]$, i.e., there is an atom $K \in \mathcal{S}[C]$ where $E \sqsubseteq^{\emptyset} K$. Then $D \sqsubseteq^{\mathcal{T}} K$ holds and thus K is in $\mathcal{S}[C]\restriction_D$.

Now since $\mathcal{T}$ contains $C \sqsubseteq D$, the pair $(x_{D,\mathcal{S}[C]\restriction_D}, x_{C,\mathcal{S}[C]})$ is contained in the simulation $\mathfrak{S}$ in Lemma VIII. It follows that $\exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{C,\mathcal{S}[C]}) \sqsubseteq^{\emptyset} \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{D,\mathcal{S}[C]\restriction_D})$ [20, Proposition 3.4.2], which implies the subsumption $\sqsubseteq^{\emptyset}$ in the claimed equivalence. Since $\mathcal{S}[C]\restriction_D$ is tolerated by $x_{C,\mathcal{S}[C]}$, the converse subsumption follows from Lemma XIII. □

Thirdly, when computing the product of D and $\exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{D,\mathcal{S}[C \sqsubseteq D]})$ we do not need to consider all objects $x_{E,\mathcal{K}}$ that are reachable from $x_{D,\mathcal{S}[C \sqsubseteq D]}$ in $\mathcal{J}_{\mathcal{S}}$, but only those where E is a filler of an existential restriction that occurs in D.

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

$$E{\restriction_{\mathcal{K}}} \coloneqq \bigsqcap\{\, A \mid A \in \mathsf{Conj}(E) \text{ and } A \notin \mathcal{K} \,\} \\ \sqcap \bigsqcap \left\{ \exists r.F{\restriction_{\mathcal{L}}} \;\middle|\; \begin{array}{l} \exists r.F \in \mathsf{Conj}(E), \text{ and } \mathcal{L} \text{ is a } \leq^{\emptyset}\text{-minimal mod. type} \\ \text{for } x_F \text{ that respects } \mathcal{S} \text{ and where } \mathsf{Succ}(\mathcal{K},r,x_F) \leq^{\emptyset} \mathcal{L} \end{array} \right\}$$

Lemma XV. $E \vee \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{E,\mathcal{K}}) \equiv^{\emptyset} E{\restriction_{\mathcal{K}}}$

Proof. We show the claim by induction on E.

Suppose a concept name A in $\mathsf{Conj}(E{\restriction_{\mathcal{K}}})$, i.e., we have $A \in \mathsf{Conj}(E)$ and $A \notin \mathcal{K}$. Since the former implies $E \sqsubseteq^{\mathcal{T}} A$, we infer that $x_{E,\mathcal{K}} \in A^{\mathcal{J}_{\mathcal{S}}}$ and thus that A is a top-level conjunct of $E \vee \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{E,\mathcal{K}})$.

Conversely, let A be a concept name in the top-level conjunction of $E \vee \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{E,\mathcal{K}})$, i.e., $A \in \mathsf{Conj}(E)$, $E \sqsubseteq^{\mathcal{T}} A$, and $A \notin \mathcal{K}$. It follows that A is also a top-level conjunct of $E{\restriction_{\mathcal{K}}}$.

Next, we are concerned with the existential restrictions. Let $\exists r.F{\restriction_{\mathcal{L}}}$ be a top-level conjunct of $E{\restriction_{\mathcal{K}}}$, which means that $\exists r.F \in \mathsf{Conj}(E)$, and $\mathcal{L}$ is a $\leq^{\emptyset}$-minimal modification type for x_F that respects $\mathcal{S}$ and satisfies $\mathsf{Succ}(\mathcal{K},r,F) \leq^{\emptyset} \mathcal{L}$. It follows that $(x_{E,\mathcal{K}}, x_{F,\mathcal{L}}) \in r^{\mathcal{J}_{\mathcal{S}}}$ and so the existential restriction $\exists r.(F \vee \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{F,\mathcal{L}}))$ is a top-level conjunct in $E \vee \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{E,\mathcal{K}})$. The induction hypothesis yields that the filler $F \vee \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{F,\mathcal{L}})$ is equivalent to $F{\restriction_{\mathcal{L}}}$. We conclude that $E \vee \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{E,\mathcal{K}})$ is subsumed by $\exists r.F{\restriction_{\mathcal{L}}}$.

Last, assume that $\exists r.(F \vee \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{G,\mathcal{M}}))$ is a top-level conjunct of $E \vee \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{E,\mathcal{K}})$, i.e., it holds that $\exists r.F \in \mathsf{Conj}(E)$, $E \sqsubseteq^{\mathcal{T}} \exists r.G$, and $\mathsf{Succ}(\mathcal{K},r,G) \leq^{\emptyset} \mathcal{M}$.

We will construct a modification type $\mathcal{L}'$ for x_F that respects $\mathcal{S}$, is tolerated by $x_{G,\mathcal{M}}$, and satisfies $\mathsf{Succ}(\mathcal{K},r,F) \leq^{\emptyset} \mathcal{L}'$. Then there must be a $\leq^{\emptyset}$-minimal $\mathcal{S}$-respecting modification type $\mathcal{L}$ for x_F such that $\mathcal{L} \leq^{\emptyset} \mathcal{L}'$ and $\mathsf{Succ}(\mathcal{K},r,F) \leq^{\emptyset} \mathcal{L}$. Lemma XI yields that $\mathcal{L}$ is tolerated by $x_{G,\mathcal{M}}$ as well. Then Lemma XIII together with the induction hypothesis yields that $F{\restriction_{\mathcal{L}}}$ is subsumed by $F \vee \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{G,\mathcal{M}})$. Furthermore, $E{\restriction_{\mathcal{K}}}$ contains the top-level conjunct $\exists r.F{\restriction_{\mathcal{L}}}$, which implies that $E{\restriction_{\mathcal{K}}}$ is subsumed by $\exists r.(F \vee \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{G,\mathcal{M}}))$.

Initialize $\mathcal{L}' \coloneqq \emptyset$. We first add enough atoms to $\mathcal{L}'$ such that it covers $\mathsf{Succ}(\mathcal{K},r,F)$. Suppose that $H \in \mathsf{Succ}(\mathcal{K},r,F)$, i.e., $\exists r.H \in \mathcal{K}$ and $F \sqsubseteq^{\mathcal{T}} H$.

- If $G \sqsubseteq^{\mathcal{T}} H$, then $H \in \mathsf{Succ}(\mathcal{K},r,G)$. Since the latter is covered by $\mathcal{M}$, there is an atom M in $\mathcal{M}$ such that $H \sqsubseteq^{\emptyset} M$. Replace $\mathcal{L}'$ by $\mathsf{Max}(\mathcal{L}' \cup \{M\})$.
- In the remaining case it holds that $G \not\sqsubseteq^{\mathcal{T}} H$. Then there must be an atom $H' \in \mathsf{Conj}(H)$ such that $G \not\sqsubseteq^{\mathcal{T}} H'$. Replace $\mathcal{L}'$ by $\mathsf{Max}(\mathcal{L}' \cup \{H'\})$.

The so obtained set $\mathcal{L}'$ is already a modification type for x_F that is tolerated by $x_{G,\mathcal{M}}$, but it might not respect $\mathcal{S}$. As shown in Lemma XII, we can extend $\mathcal{L}'$ by means of the Extension Rule to an $\mathcal{S}$-respecting modification type. ☐

The following characterization of the conclusions in the induced GC-repair $\mathsf{rep}_{\mathsf{GC}}(\mathcal{T},\mathcal{S})$ follows from Lemmas VII, IX, XIV, and XV, and it is the main result of this section.

Proposition 15. *Given a repair seed $\mathcal{S}$, it holds that $D \vee C^{\mathcal{J}_\mathcal{S}\mathcal{J}_\mathcal{S}} \equiv^\emptyset D{\upharpoonright}_{\mathcal{S}[C \sqsubseteq D]}$ for each CI $C \sqsubseteq D$ in $\mathcal{T}$, and thus the induced GC-repair $\mathsf{rep}_{\mathsf{GC}}(\mathcal{T}, \mathcal{S})$ is equivalent to the TBox $\{\, C \sqsubseteq D{\upharpoonright}_{\mathcal{S}[C \sqsubseteq D]} \mid C \sqsubseteq D \in \mathcal{T} \,\}$.*

It remains an open question whether it is tractable to decide if $\mathsf{rep}_{\mathsf{GC}}(\mathcal{T}, \mathcal{S}_1)$ is a GC-weakening of $\mathsf{rep}_{\mathsf{GC}}(\mathcal{T}, \mathcal{S}_2)$, i.e., whether this can be decided in polynomial time w.r.t. the size of the input TBox $\mathcal{T}$ and the repair request $\mathcal{P}$. However, under the assumption that $\mathcal{T}$ does not contain multiple CIs with the same premise, we can recursively decide if $\mathsf{rep}_{\mathsf{GC}}(\mathcal{T}, \mathcal{S}_1) \succeq_{\mathsf{GC}} \mathsf{rep}_{\mathsf{GC}}(\mathcal{T}, \mathcal{S}_2)$ holds without always constructing the whole induced GC-repairs. By the above proposition it suffices to check if $D{\upharpoonright}_{\mathcal{S}_1[C \sqsubseteq D]} \sqsubseteq^\emptyset D{\upharpoonright}_{\mathcal{S}_2[C \sqsubseteq D]}$ for each $C \sqsubseteq D$ in $\mathcal{T}$. Using the recursive characterization of subsumption in $\mathcal{EL}$ (w.r.t. an empty TBox), the two involved concepts need to be constructed only up to the first clash.

3.4 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 $\mathcal{T} \coloneqq \{A \sqsubseteq B,\ C \sqsubseteq \exists r.(A \sqcap B)\}$ and the repair request $\mathcal{P} \coloneqq \{C \sqsubseteq \exists r.B\}$, there are two optimal GC-repairs: $\mathcal{U}_1 \coloneqq \{A \sqsubseteq B,\ C \sqsubseteq \exists r.\top\}$, induced by the seed $\mathcal{S}_1 \coloneqq \{A \sqsubseteq B\}$, and $\mathcal{U}_2 \coloneqq \{A \sqsubseteq \top,\ C \sqsubseteq \exists r.A\}$, induced by $\mathcal{S}_2 \coloneqq \emptyset$. 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 $F \in \mathsf{Atoms}(\mathcal{P}, \mathcal{T})$ such that $D \sqsubseteq^\emptyset F$. We definitely sometimes need CIs $C \sqsubseteq F$ where $C \sqsubseteq^\mathcal{T} F$, 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 $\mathcal{T} \coloneqq \{A \sqsubseteq \exists r.\exists r.(B \sqcap C),\ \exists r.B \sqsubseteq B\}$ and the repair request $\mathcal{P} \coloneqq \{A \sqsubseteq \exists r.\exists r.C\}$. 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 $\mathcal{T} \models_{\mathsf{FP}} \mathcal{U}$, if $\mathsf{Prem}(\mathcal{T}) = \mathsf{Prem}(\mathcal{U})$ and $\mathcal{T} \models \mathcal{U}$.*

$\mathcal{T} \succeq_{\mathsf{GC}} \mathcal{U}$ implies $\mathcal{T} \models_{\mathsf{FP}} \mathcal{U}$ and the latter implies $\mathcal{T} \models \mathcal{U}$, but the converse implications need not hold. This means that the relation $\models_{\mathsf{FP}}$ is between $\succeq_{\mathsf{GC}}$ 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 $\mathcal{V} \models_{\mathsf{FP}} \mathcal{U}$ and $\mathcal{U} \not\models_{\mathsf{FP}} \mathcal{V}$.*

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*

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

Recall that each conclusion $C^{\mathcal{J}_{\mathcal{S}}\mathcal{J}_{\mathcal{S}}}$ is equivalent to the $\mathcal{EL}_{\mathsf{si}}$ concept description $\exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{C,\mathcal{S}[C]})$, where $\mathcal{S}[C]$ is the largest modification type for x_C that respects $\mathcal{S}$ and does not cover $\{C\}$, i.e., $\mathcal{S}[C] \coloneqq \mathsf{Max}\{\, K \mid K \in \mathsf{Atoms}(\mathcal{P}, \mathcal{T}),\ C \not\sqsubseteq^{\mathcal{S}} K, \text{ and } C \sqsubseteq^{\mathcal{T}} K \,\}$. Analogously to the GC-repairs, every TBox $\mathsf{rep}_{\mathsf{FP}}(\mathcal{T}, \mathcal{S})$ is an FP-repair and each FP-repair is FP-entailed by $\mathsf{rep}_{\mathsf{FP}}(\mathcal{T}, \mathcal{S})$ for some repair seed $\mathcal{S}$.

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

Proof. Recall from Section 2 that a CI $C \sqsubseteq D$ is satisfied in $\mathcal{J}_{\mathcal{S}}$ iff $C^{\mathcal{J}_{\mathcal{S}}\mathcal{J}_{\mathcal{S}}} \sqsubseteq^{\emptyset} D$. It follows that the induced countermodel $\mathcal{J}_{\mathcal{S}}$ is a model of $\mathsf{rep}_{\mathsf{FP}}(\mathcal{T}, \mathcal{S})$. According to Proposition 7, $\mathcal{J}_{\mathcal{S}}$ does not satisfy any CI in $\mathcal{P}$. It follows that $\mathsf{rep}_{\mathsf{FP}}(\mathcal{T}, \mathcal{S})$ complies with $\mathcal{P}$.

It remains to show that $\mathsf{rep}_{\mathsf{FP}}(\mathcal{T}, \mathcal{S})$ is FP-entailed by $\mathcal{T}$. Definition 20 implies that each premise in $\mathsf{rep}_{\mathsf{FP}}(\mathcal{T}, \mathcal{S})$ is a premise in $\mathcal{T}$ too. Now consider a CI $C \sqsubseteq C^{\mathcal{J}_{\mathcal{S}}\mathcal{J}_{\mathcal{S}}}$ in $\mathsf{rep}_{\mathsf{FP}}(\mathcal{T}, \mathcal{S})$. We must prove that it is entailed by $\mathcal{T}$. Recall from Lemma IX that $C^{\mathcal{J}_{\mathcal{S}}\mathcal{J}_{\mathcal{S}}} \equiv^{\emptyset} \exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{C,\mathcal{S}[C]})$. Let $\mathcal{J}$ be a model of $\mathcal{T}$ such that $y \in C^{\mathcal{J}}$. Lemma II yields that there is a simulation from $\mathcal{I}$ to $\mathcal{J}$ containing (x_C, y). Furthermore, it is easy to see that the relation $\{\, (x_{D,\mathcal{K}}, x_D) \mid x_{D,\mathcal{K}} \in \mathsf{Dom}(\mathcal{J}_{\mathcal{S}}) \,\}$ is a simulation from $\mathcal{J}_{\mathcal{S}}$ to $\mathcal{I}$ that contains $(x_{C,\mathcal{S}[C]}, x_C)$. Composing these two simulations yields one from $\mathcal{J}_{\mathcal{S}}$ to $\mathcal{J}$ that contains $(x_{C,\mathcal{S}[C]}, y)$, i.e., $y \in (\exists^{\mathsf{sim}}(\mathcal{J}_{\mathcal{S}}, x_{C,\mathcal{S}[C]}))^{\mathcal{J}}$ and we are done. □

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

Proof. Consider an FP-repair $\mathcal{U}$ of $\mathcal{T}$ for $\mathcal{P}$, i.e., every premise in $\mathcal{U}$ is also a premise in $\mathcal{T}$. As in the proof of Proposition 11, we can construct a repair seed $\mathcal{S}^*_{\mathcal{U}}$ such that its induced countermodel $\mathcal{J}_{\mathcal{S}^*_{\mathcal{U}}}$ is a model of $\mathcal{U}$. It follows that $C^{\mathcal{J}_{\mathcal{S}^*_{\mathcal{U}}}\mathcal{J}_{\mathcal{S}^*_{\mathcal{U}}}} \sqsubseteq^{\emptyset} D$ for each CI $C \sqsubseteq D \in \mathcal{U}$, and furthermore $\mathsf{rep}_{\mathsf{FP}}(\mathcal{T}, \mathcal{S})$ contains the CI $C \sqsubseteq C^{\mathcal{J}_{\mathcal{S}^*_{\mathcal{U}}}\mathcal{J}_{\mathcal{S}^*_{\mathcal{U}}}}$. We conclude that $\mathsf{rep}_{\mathsf{FP}}(\mathcal{T}, \mathcal{S})$ entails each CI in $\mathcal{U}$. □

We obtain the following main result of this section. Its proof is analogous to Theorem 12, but uses the argument that entailment between $\mathcal{EL}_{\mathsf{si}}$ TBoxes can be decided in polynomial time [25].

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. We need the following lemma to prove this.

Lemma XVI. *Let $\mathcal{S}$ be a repair seed. The induced countermodel $\mathcal{J}_{\mathcal{S}}$ is a model of $\mathcal{S}$.*

Proof. Consider a concept inclusion $C \sqsubseteq F$ in the seed $\mathcal{S}$, which specifically means that $C \in \mathsf{Prem}(\mathcal{T})$, $F \in \mathsf{Atoms}(\mathcal{P}, \mathcal{T})$ and $C \sqsubseteq^{\mathcal{T}} F$, cf. Definition 4. Further let $x_{E,\mathcal{K}} \in C^{\mathcal{J}_{\mathcal{S}}}$. By Lemma VI it follows that $x_E \in C^{\mathcal{I}}$ and $\{C\} \not\leq^{\emptyset} \mathcal{K}$.

We show that $\{F\} \not\leq^{\mathcal{S}} \mathcal{K}$. Assume that this was not the case. Then $\{C\} \leq^{\mathcal{S}} \mathcal{K}$ would hold. Since $\mathcal{K}$ respects $\mathcal{S}$, it would follow that $\{C\} \leq^{\emptyset} \mathcal{K}$ — a contradiction.

Since $\mathcal{I}$ is a model of $\mathcal{T}$ by Lemma III, we infer that $x_E \in F^{\mathcal{I}}$. An application of Lemma VI yields that $x_{E,\mathcal{K}} \in F^{\mathcal{J}_{\mathcal{S}}}$. □

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

Proof. Consider a concept inclusion $C \sqsubseteq F$ in the seed $\mathcal{S}$. Recall from Definition 4 that C is a premise in $\mathcal{T}$. According to Lemma XVI, the induced countermodel $\mathcal{J}_{\mathcal{S}}$ is a model of $\mathcal{S}$, which implies $C^{\mathcal{J}_{\mathcal{S}}\mathcal{J}_{\mathcal{S}}} \sqsubseteq^{\emptyset} F$. Since the FP-repair $\mathsf{rep}_{\mathsf{FP}}(\mathcal{T}, \mathcal{S})$ contains the CI $C \sqsubseteq C^{\mathcal{J}_{\mathcal{S}}\mathcal{J}_{\mathcal{S}}}$, we conclude that $\mathsf{rep}_{\mathsf{FP}}(\mathcal{T}, \mathcal{S})$ entails $C \sqsubseteq F$. □

Contrary to the GC-repairs, not every FP-repair is an $\mathcal{EL}$ TBox but might require cyclic $\mathcal{EL}_{\mathsf{si}}$ concept descriptions [25] 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

$$A \sqsubseteq \exists^{\mathsf{sim}}(\ \rightarrow \overset{A}{\bigcirc} \xrightarrow{r} \bigcirc \xrightarrow{r} \overset{A}{\bigcirc} \circlearrowleft r\ ).$$

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 [26, 30] 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 $\mathsf{rep}_{\mathsf{FP}}(\mathcal{T}, \mathcal{S})$ then only has acyclic $\mathcal{EL}_{\mathsf{si}}$ concept descriptions as conclusions and these can be rewritten into $\mathcal{EL}$ concept descriptions.

Prototypical Implementation

The two repair approaches from the previous sections have been implemented in a prototype, available at `https://github.com/francesco-kriegel/right-repairs-of-el-tboxes`. The needed repair seed is obtained by interaction with the user, who must specify which of the polynomially many CIs of the form $C \sqsubseteq F$ as per Definition 4 are valid — all confirmed CIs constitute the repair seed. However, the user need not check each such CI. Firstly, there is no need to ask whether a tautology is valid. Secondly, if a CI follows from previously confirmed CIs, then it must not be rejected and so there is no need to ask for it. Thirdly, since the goal is to build a repair seed, a CI must be rejected if it together with all previously accepted CIs entails one of the unwanted consequences in $\mathcal{P}$. Only the remaining CIs need to be decided by the user.

Currently, the CIs are presented to the user in the following order: $C_1 \sqsubseteq F_1$ comes before $C_2 \sqsubseteq F_2$ if C_1 subsumes a subconcept C_2 (since then the former could potentially be used to entail the latter). Other orders could also be suitable. For instance, the implementation could be changed such that the user decides $C_1 \sqsubseteq F_1$ before $C_2 \sqsubseteq F_2$ if the latter does not follow from all previously accepted CIs but it would follow if also the former would be accepted.

5 Complexity of Maximally Strong $\succ^{\mathsf{sub}}$-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 $\mathcal{O}_{\mathsf{s}}$ and a refutable part $\mathcal{O}_{\mathsf{r}}$ as well as an axiom α, the unwanted consequence, that follows from $\mathcal{O}$ but not already from $\mathcal{O}_{\mathsf{s}}$. A *repair* is an ontology $\mathcal{O}'$ such that $\mathcal{O} \models \mathcal{O}'$ but $\mathcal{O}_{\mathsf{s}} \cup \mathcal{O}' \not\models \alpha$. 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.[5]

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

Specifically, a *justification* for α is a minimal subset $J \subseteq \mathcal{O}_\mathsf{r}$ such that $\mathcal{O}_\mathsf{s} \cup J \models \alpha$. 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 $\mathcal{O}_\mathsf{s} \cup (J \setminus \{\beta\}) \cup \{\gamma\} \not\models \alpha$, but $\mathcal{O}_\mathsf{s} \cup (J \setminus \{\beta\}) \cup \{\delta\} \models \alpha$ 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* $\succ^\mathsf{syn}$ on $\mathcal{EL}$ CIs removes subconcepts from the conclusions. Maximally strong $\succ^\mathsf{syn}$-weakenings always exist in all directions,[6] 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* $\succ^\mathsf{sub}$ replaces conclusions of $\mathcal{EL}$ CIs by more general concepts, i.e., $C \sqsubseteq D \succ^\mathsf{sub} C' \sqsubseteq D'$ if $C = C'$, $D \sqsubset^\emptyset D'$, and $C' \sqsubseteq D' \not\models C \sqsubseteq D$. It has only been known that maximally strong $\succ^\mathsf{sub}$-weakenings always exist in all directions,[6] all of them can effectively be computed, and recognizing them is coNP-hard. As a side result from Section 3, we obtain that the former can actually be done in exponential time if the unwanted consequence α is a CI. We further show that the latter is also an upper bound, and that a single maximally strong $\succ^\mathsf{sub}$-weakening cannot be computed in polynomial time.

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

Proof. Fix a justification J for the unwanted consequence α and let $\beta \in J$. Recall that then α follows from $\mathcal{O}_\mathsf{s} \cup J$ but, due to minimality, not from $\mathcal{O}_\mathsf{s} \cup (J \setminus \{\beta\})$. Since in $\mathcal{EL}$ a CI follows from an ontology iff it already follows from the CIs in the ontology, J consists of CIs only. Let the TBox $\mathcal{T}$ consist of all concept inclusions in the static part $\mathcal{O}_\mathsf{s}$ or in the justification J, and further define the repair request as $\mathcal{P} := \{\alpha\}$. We then have that $\mathcal{T} \models \alpha$ but $\mathcal{T} \setminus \{\beta\} \not\models \alpha$. First of all, if γ is a maximally strong $\succ^\mathsf{sub}$-weakening of β in J, then $(\mathcal{T} \setminus \{\beta\}) \cup \{\gamma\}$ is a GC-repair of $\mathcal{T}$ for $\mathcal{P}$. It must be an optimal one, since otherwise γ would not be maximally strong. We conclude that each maximally strong $\succ^\mathsf{sub}$-weakening of β in J corresponds to an optimal GC-repair.

The repair seeds that induce these optimal GC-repairs have a specific form. Assume that $\mathcal{U} := (\mathcal{T} \setminus \{\beta\}) \cup \{\gamma\}$ where γ is a maximally strong $\succ^\mathsf{sub}$-weakening of β in J. According to the inductive construction of the repair seed $\mathcal{S}^*_\mathcal{U}$ in the proof of Proposition 11, the initial set $\mathcal{S}^1_\mathcal{U} \subseteq \mathcal{S}^*_\mathcal{U}$ contains the concept inclusions $C \sqsubseteq F$ for all $C \sqsubseteq D \in \mathcal{T} \setminus \{\beta\}$ and all $F \in \mathsf{Conj}(D)$, possibly among others. It follows that the final seed $\mathcal{S}^*_\mathcal{U}$ entails $\mathcal{T} \setminus \{\beta\}$. In addition, $\mathcal{S}^*_\mathcal{U}$ does not entail α.

[6] 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 $\mathcal{O}_\mathsf{s} \cup (J \setminus \{\beta\}) \cup \{\gamma\} \not\models \alpha$.

We conclude that, in order to compute the maximally strong $\succ^{\mathsf{sub}}$-weakenings of β in J, it suffices to consider only those repair seeds $\mathcal{S}$ that entail $\mathcal{T} \setminus \{\beta\}$ but do not entail α. It is easy to see that then the modification type $\mathcal{S}[C \sqsubseteq D]$ is empty for each $C \sqsubseteq D \in \mathcal{T} \setminus \{\beta\}$, and thus the induced repair $\mathsf{rep}_{\mathsf{GC}}(\mathcal{T}, \mathcal{S})$ still contains, possibly up to equivalence, all CIs in $\mathcal{T} \setminus \{\beta\}$, cf. Proposition 15. It further contains one additional CI with premise C_β where $\beta =: C_\beta \sqsubseteq D_\beta$, which is a $\succ^{\mathsf{sub}}$-weakening of β and which we denote by $\gamma_{\mathcal{S}}$. Due to the above correspondence between maximally strong $\succ^{\mathsf{sub}}$-weakenings and the special repair seeds, we can finally find every maximally strong $\succ^{\mathsf{sub}}$-weakening of β among these CIs $\gamma_{\mathcal{S}}$. Since there are at most exponentially many such seeds $\mathcal{S}$, there at most exponentially many such $\succ^{\mathsf{sub}}$-weakenings $\gamma_{\mathcal{S}}$. Since entailment in $\mathcal{EL}$ can be decided in polynomial time [3], we can identify the maximally strong ones among them in exponential time, and so the claim follows. □

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

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

Finally, recognizing maximally strong $\succ^{\mathsf{sub}}$-weakenings is also in $\mathsf{co}\,\mathsf{NP}$.

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

Proof. Hardness was already shown [8]. We now turn our attention to proving containment. We first cite a special case of the definition of a most specific consequence [21, Definition 3]: given an $\mathcal{EL}$ TBox $\mathcal{T}$ and an $\mathcal{EL}$ concept description C, an $\mathcal{EL}_{\mathsf{si}}$ concept description D is called *most specific consequence* of C w.r.t. $\mathcal{T}$ if, firstly, $C \sqsubseteq^{\mathcal{T}} D$ and, secondly, $C \sqsubseteq^{\mathcal{T}} E$ implies $D \sqsubseteq^{\emptyset} E$ for each $\mathcal{EL}_{\mathsf{si}}$ concept description E. Furthermore, the most specific consequence of C w.r.t. $\mathcal{T}$ always exists and is equivalent to the $\mathcal{EL}_{\mathsf{si}}$ concept description $\exists^{\mathsf{sim}}(\mathcal{I}_{C,\mathcal{T}}, x_C)$ [21, Proposition 6].[7] Up to equivalence, we will denote it as $C^{\mathcal{T}}$.

Assume that J is a justification for α, i.e., $\mathcal{O}_{\mathsf{s}} \cup J \models \alpha$ but no proper subset of J satisfies this, and let $\beta \coloneqq C \sqsubseteq D$ be a CI in J. Further consider a CI $\gamma \coloneqq C \sqsubseteq E$. It is easy to check in polynomial time [3] whether $C \sqsubseteq E$ is a $\succ^{\mathsf{sub}}$-weakening of $C \sqsubseteq D$: one only needs to check if $D \sqsubseteq^{\emptyset} E$ and if $\mathcal{O}_{\mathsf{s}} \cup (J \setminus \{C \sqsubseteq D\}) \cup \{C \sqsubseteq E\} \not\models \alpha$. Assume in the following that this test succeeds. It then remains to check if $C \sqsubseteq E$ is maximally strong.

[7] $\mathcal{I}_{C,\mathcal{T}}$ is the *canonical model* of C and $\mathcal{T}$ [27], which can alternatively be defined like the canonical model $\mathcal{I}$ in Section 3.1 but using the domain $\{\, x_D \mid D \in \mathsf{Sub}(C, \mathcal{T}) \,\}$.

Next, we define a normal form of $C \sqsubseteq E$. Therefore, consider the $\mathcal{EL}$ concept description $E^{\downarrow} := D \vee C^{\{C \sqsubseteq E\}}$. Recall from Section 2 that $E^{\downarrow}$ must be an $\mathcal{EL}$ concept description since D is one and the most specific consequence $C^{\{C \sqsubseteq E\}}$ is an $\mathcal{EL}_{\mathsf{si}}$ concept description. This concept description $E^{\downarrow}$ can be computed in polynomial time for the following reasons. The most specific consequence $C^{\{C \sqsubseteq E\}}$ is equivalent to the $\mathcal{EL}_{\mathsf{si}}$ concept description $\exists^{\mathsf{sim}}(\mathcal{I}_{C,\{C \sqsubseteq E\}}, x_C)$, which can be computed in polynomial time since the canonical model $\mathcal{I}_{C,\{C \sqsubseteq E\}}$ can be computed in polynomial time. The least common subsumer of D and $\exists^{\mathsf{sim}}(\mathcal{I}_{C,\{C \sqsubseteq E\}}, x_C)$ can be obtained as the product of the syntax tree of D and the interpretation $\mathcal{I}_{C,\{C \sqsubseteq E\}}$, rooted at x_C. This product can clearly be computed in polynomial time.

We show the following three claims.

1. $D \sqsubseteq^{\emptyset} E^{\downarrow} \sqsubseteq^{\emptyset} E$

   By definition, we have $D \sqsubseteq^{\emptyset} E^{\downarrow}$. Since $C \sqsubseteq E$ is a $\succ^{\mathsf{sub}}$-weakening of $C \sqsubseteq D$, we have $D \sqsubseteq^{\emptyset} E$. It is further easy to see that $C^{\{C \sqsubseteq E\}} \sqsubseteq^{\emptyset} E$. It follows that the least common subsumer of D and $C^{\{C \sqsubseteq E\}}$, which is $E^{\downarrow}$, is subsumed by E.
2. $C \sqsubseteq E^{\downarrow}$ and $C \sqsubseteq E$ are equivalent.

   Clearly, $E^{\downarrow} \sqsubseteq^{\emptyset} E$ implies that $C \sqsubseteq E^{\downarrow}$ entails $C \sqsubseteq E$. Again by definition, it holds that $C^{\{C \sqsubseteq E\}} \sqsubseteq^{\emptyset} E^{\downarrow}$, and thus $C \sqsubseteq C^{\{C \sqsubseteq E\}}$ entails $C \sqsubseteq E^{\downarrow}$. Furthermore, the TBoxes $\{C \sqsubseteq E\}$ and $\{C \sqsubseteq C^{\{C \sqsubseteq E\}}\}$ are equivalent [21, Lemma 20]. We infer that also $C \sqsubseteq E$ entails $C \sqsubseteq E^{\downarrow}$.
3. $E^{\downarrow}$ is most specific, i.e., for each $\mathcal{EL}$ concept description E', if $D \sqsubseteq^{\emptyset} E' \sqsubseteq^{\emptyset} E$, and $C \sqsubseteq E'$ and $C \sqsubseteq E$ are equivalent, then $E^{\downarrow} \sqsubseteq^{\emptyset} E'$.

   Assume that $D \sqsubseteq^{\emptyset} E' \sqsubseteq^{\emptyset} E$ and that $C \sqsubseteq E'$ and $C \sqsubseteq E$ are equivalent. We obtain $E^{\downarrow} = D \vee C^{\{C \sqsubseteq E\}} \equiv^{\emptyset} D \vee C^{\{C \sqsubseteq E'\}} \sqsubseteq^{\emptyset} D \vee E' \equiv^{\emptyset} E'$.

Thus, $C \sqsubseteq E^{\downarrow}$ is also a $\succ^{\mathsf{sub}}$-weakening of $C \sqsubseteq D$ that is equivalent to $C \sqsubseteq E$.

Before we continue with the proof, we cite a further definition. Given $\mathcal{EL}$ concept descriptions G and L, we call L a *lower neighbor* of G, written $L \prec^{\emptyset} G$, if $L \sqsubset^{\emptyset} G$ and there is no $\mathcal{EL}$ concept description M with $L \sqsubset^{\emptyset} M \sqsubset^{\emptyset} G$ [22]. Up to equivalence, each lower neighbor has polynomial size. If $H \sqsubset^{\emptyset} G$, then there is a lower neighbor L of G such that $H \sqsubseteq^{\emptyset} L$.

Last, we show the following claim: $\gamma = C \sqsubseteq E$ is no maximally strong $\succ^{\mathsf{sub}}$-weakening of $\beta = C \sqsubseteq D$ in J iff there is a concept description L such that $D \sqsubseteq^{\emptyset} L \prec^{\emptyset} E^{\downarrow}$ and $\mathcal{O}_{\mathsf{s}} \cup (J \setminus \{C \sqsubseteq D\}) \cup \{C \sqsubseteq L\} \not\models \alpha$.

Regarding the if direction, let L be an $\mathcal{EL}$ concept description such that $D \sqsubseteq^{\emptyset} L \prec^{\emptyset} E^{\downarrow}$ and $\mathcal{O}_{\mathsf{s}} \cup (J \setminus \{C \sqsubseteq D\}) \cup \{C \sqsubseteq L\} \not\models \alpha$. The latter together with $D \sqsubseteq^{\emptyset} L$ implies that $C \sqsubseteq L$ is a $\succ^{\mathsf{sub}}$-weakening of $C \sqsubseteq D$ in J. Since L is a lower neighbor of $E^{\downarrow}$, we have $L \sqsubset^{\emptyset} E^{\downarrow}$. Claim 3 implies that $C \sqsubseteq E$ and $C \sqsubseteq L$ are not equivalent, and thus $C \sqsubseteq L$ is stronger than $C \sqsubseteq E$ w.r.t. $\succ^{\mathsf{sub}}$. We conclude that $C \sqsubseteq E$ is no maximally strong $\succ^{\mathsf{sub}}$-weakening of $C \sqsubseteq D$ in J.

The only-if direction remains. For this purpose, consider a $\succ^{\mathsf{sub}}$-weakening $C \sqsubseteq F$ of $C \sqsubseteq D$ in J that is $\succ^{\mathsf{sub}}$-stronger than $C \sqsubseteq E$. Then $D \sqsubseteq^{\emptyset} F \sqsubseteq^{\emptyset} E$, and $C \sqsubseteq F$ entails $C \sqsubseteq E$ but they are not equivalent. Since $C \sqsubseteq F$ entails $C \sqsubseteq E$ and

$C \sqsubseteq E$ is equivalent to $C \sqsubseteq E^{\downarrow}$, it follows that $C^{\{C \sqsubseteq F\}} \sqsubseteq^{\emptyset} E^{\downarrow}$ [21, Definition 3]. Since also $D \sqsubseteq^{\emptyset} E^{\downarrow}$ by Claim 1, we infer that $E^{\downarrow}$ is a common subsumer of D and $C^{\{C \sqsubseteq F\}}$. Defining $F^{\downarrow}$ as the least common subsumer $D \vee C^{\{C \sqsubseteq F\}}$ like above yields $F^{\downarrow} \sqsubseteq^{\emptyset} E^{\downarrow}$.

Claim 2 yields that $C \sqsubseteq F$ and $C \sqsubseteq F^{\downarrow}$ are equivalent, as are $C \sqsubseteq E$ and $C \sqsubseteq E^{\downarrow}$. Now $F^{\downarrow} \sqsubseteq^{\emptyset} E^{\downarrow}$ implies that $C \sqsubseteq F^{\downarrow}$ entails $C \sqsubseteq E^{\downarrow}$. If $E^{\downarrow} \sqsubseteq^{\emptyset} F^{\downarrow}$, then $C \sqsubseteq E^{\downarrow}$ would entail $C \sqsubseteq F^{\downarrow}$ and thus $C \sqsubseteq F$ and $C \sqsubseteq E$ would be equivalent, a contradiction. We conclude that $F^{\downarrow} \sqsubset^{\emptyset} E^{\downarrow}$. So there exists a lower neighbor L where $F^{\downarrow} \sqsubseteq^{\emptyset} L \prec^{\emptyset} E^{\downarrow}$ [22]. Recall that $C \sqsubseteq F$ is a $\succ^{\mathsf{sub}}$-weakening of $C \sqsubseteq D$ in J and that $C \sqsubseteq F$ and $C \sqsubseteq F^{\downarrow}$ are equivalent by Claim 2, i.e., we have $\mathcal{O}_{\mathsf{s}} \cup (J \setminus \{C \sqsubseteq D\}) \cup \{C \sqsubseteq F^{\downarrow}\} \not\models \alpha$. Since $C \sqsubseteq F^{\downarrow}$ entails the CI $C \sqsubseteq L$, it follows that $\mathcal{O}_{\mathsf{s}} \cup (J \setminus \{C \sqsubseteq D\}) \cup \{C \sqsubseteq L\} \not\models \alpha$.

Consequently, to check if $C \sqsubseteq E$ is maximally strong we only need to determine whether there is some lower neighbor L of $E^{\downarrow}$ such that $D \sqsubseteq^{\emptyset} L$ and $\mathcal{O}_{\mathsf{s}} \cup (J \setminus \{C \sqsubseteq D\}) \cup \{C \sqsubseteq L\} \not\models \alpha$. If so, then L is a certificate for $C \sqsubseteq E$ being not maximally strong. Otherwise, if there is no such L, then $C \sqsubseteq E$ is maximally strong. Recall that, up to equivalence, each lower neighbor has polynomial size [22]. This means that we can guess a concept L of polynomial size and then check if it subsumes D, is a lower neighbor of $E^{\downarrow}$, and satisfies $\mathcal{O}_{\mathsf{s}} \cup (J \setminus \{C \sqsubseteq D\}) \cup \{C \sqsubseteq L\} \not\models \alpha$. So the containment in coNP follows. $\square$

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 $\bot$, nominals $\{a\}$, inverse roles r^-, and role inclusions $R_1 \circ \cdots \circ R_n \sqsubseteq S$. Ideas from the latest extension of quantified ABox repairs to the DL $\mathcal{ELROI}(\bot)$ 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., Dachselt, 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: Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning, KR 2022, Haifa, Israel, July 31 – August 5, 2022 (2022), to appear.
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., Morawska, B.: Unification in the description logic $\mathcal{EL}$. Log. Methods Comput. Sci. **6**(3) (2010), http://arxiv.org/abs/1006.2289
11. Baader, F., Peñaloza, R.: Axiom pinpointing in general tableaux. J. Log. Comput. **20**(1), 5–34 (2010), https://doi.org/10.1093/logcom/exn058
12. 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
13. 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
14. Cóbe, 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
15. 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

16. 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`
17. Horrocks, I., Kutz, O., Sattler, U.: The even more irresistible $\mathcal{SROIQ}$. 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`
18. 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`
19. Kazakov, Y.: $\mathcal{RIQ}$ and $\mathcal{SROIQ}$ are harder than $\mathcal{SHOIQ}$. 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`
20. Kriegel, F.: Constructing and Extending Description Logic Ontologies using Methods of Formal Concept Analysis. Doctoral thesis, Technische Universität Dresden, Dresden, Germany (2019), `http://nbn-resolving.de/urn:nbn:de:bsz:14-qucosa2-360998`
21. Kriegel, F.: Most specific consequences in the description logic $\mathcal{EL}$. Discret. Appl. Math. **273**, 172–204 (2020), `https://doi.org/10.1016/j.dam.2019.01.029`
22. Kriegel, F.: Navigating the $\mathcal{EL}$ subsumption hierarchy. 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-21.pdf`
23. 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`
24. 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`
25. 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`
26. 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
27. 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
28. 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
29. 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
30. Nikitina, N., Rudolph, S.: ExpExpExplosion: uniform interpolation in general $\mathcal{EL}$ terminologies. In: De Raedt, L., Bessiere, 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
31. 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, see the erratum [15].
32. 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
33. 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
34. 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
35. 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
36. Tena Cucala, D., Cuenca Grau, B., Horrocks, I.: Pay-as-you-go consequence-based reasoning for the description logic $\mathcal{SROIQ}$. Artif. Intell. **298**, 103518 (2021), https://doi.org/10.1016/j.artint.2021.103518