Relating Optimal Repairs in Ontology Engineering with Contraction Operations in Belief Change

Franz Baader
TU Dresden, Institute of Theoretical Computer Science, Dresden, Germany
and Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI) Dresden/Leipzig, Germany
franz.baader@tu-dresden.de

ABSTRACT

The question of how a given knowledge base can be modified such that certain unwanted consequences are removed has been investigated in the area of ontology engineering under the name of repair and in the area of belief change under the name of contraction. Whereas in the former area the emphasis was more on designing and implementing concrete repair algorithms, the latter area concentrated on characterizing classes of contraction operations by certain postulates they satisfy. In the classical setting, repairs and contractions are subsets of the knowledge base that no longer have the unwanted consequence. This makes these approaches syntax-dependent and may result in removal of more consequences than necessary. To alleviate this problem, gentle repairs and pseudo-constractions have been introduced in the respective research areas, and their connections have been investigated in recent work. Optimal repairs preserve a maximal amount of consequences, but they may not always exist. We show that, if they exist, then they can be obtained by certain pseudo-contraction operations, and thus they comply with the postulates that these operations satisfy. Conversely, under certain conditions, pseudo-contractions are guaranteed to produce optimal repairs. Recently, contraction operations have also been defined for concepts rather than for whole knowledge bases. We show that there is again a close connection between such operations and optimal repairs of a restricted form of knowledge bases.

CCS Concepts

•Theory of computation → Description logics; •Computing methodologies → Ontology engineering; Nonmonotonic, default reasoning and belief revision;

Keywords

Belief Change, Ontology Repair, Description Logic

1. INTRODUCTION

Representing knowledge in a logic-based knowledge representation language allows one to draw implicit consequences from the explicitly represented knowledge. If such a consequence is deemed to be incorrect or no longer wanted for some reason, then it is often not obvious how to modify the knowledge base to get rid of this consequence. In ontology engineering, the knowledge base (also called ontology) usually defines the important notions of the application domain as background knowledge in the terminology, and then uses these notions to represent a specific application situation. Modelling errors are detected when the reasoner generates a consequence that formally follows from the knowledge base, but is incorrect in the sense that it does not hold in the application domain that is supposed to be modelled. The question is then how to repair the knowledge base such that no new consequences are added, the unwanted consequence no longer follows, and other consequences are not lost unnecessarily. The classical approaches for ontology repair consider as repairs maximal subsets of the ontology (viewed as a set of logical sentences) that do not have the unwanted consequence, and employ methods inspired by model-based diagnosis [28] to compute these sets [26, 31, 10], which are called optimal classical repairs in [8]. While these approaches preserve as many of the sentences in the ontology as possible, they need not preserve a maximal amount of consequences (see [8] as well as the examples at the end of Section 2 and the beginning of Section 4 of the present paper). To overcome this problem, more gentle repair approaches have been introduced, e.g., in [23, 32, 8], but these methods still need not produce optimal repairs, i.e., ones that preserve a maximal set of consequences. In general, such optimal repairs need not exist [8]. In the setting of repairing ABoxes of the description logic $\mathcal{EL}$ w.r.t. static $\mathcal{EL}$ TBoxes, methods for computing optimal repairs (if they exist) are available [6].

In belief change [14], one usually assumes that the knowledge base represents the beliefs of a rational agent. These beliefs may change if the agent receives new information, and the question is how this can be reflected by a change of the knowledge base. Removing (implied) information is called contraction in this setting. Instead of directly constructing contraction operations, the belief change community has formulated properties (called postulates) that should be satisfied by reasonable contraction operations, and then developed approaches for constructing contraction operations that capture exactly those contraction operations that satisfy a certain combination of postulates. This approach, which was pioneered in [1], is called the AGM approach. The original AGM approach works with *belief sets*, which are assumed to be closed under consequences. From a practical point of view, it makes more sense to work with non-deductively closed (and ideally finite) representations of belief sets, called belief bases [16, 13, 18]. Similar to classical repairs, the original approaches for belief base contraction consider subsets of the knowledge base as possible contractions. For the same reasons as for repairs, operations that preserve more consequences, called pseudo-contractions, have been introduced in the belief change literature [15, 17, 30, 24, 25].

Although contractions and classical repairs as well as pseudo-contractions and repairs tackle basically the same problems, there has until recently been little interaction between the two communities, and thus the connections between the developed approaches remained unclear. In [24, 25], this problem is addressed with an emphasis on showing connections between gentle repairs and certain pseudo-contraction approaches called partial meet and kernel pseudo-contractions. In the present paper, we concentrate on optimal repairs, both in the classical and the general sense. We show that, under certain conditions, operations that compute optimal (classical) repairs can be obtained as partial meet and kernel pseudo-contractions (contractions), and vice versa. This shows, on the one hand, that the approaches developed in ontology engineering satisfy the postulates required in belief change. On the other hand, under certain conditions the approaches developed in belief change yield optimal (classical) repairs. We instantiate our results using the setting of repairing ABoxes of the description logic $\mathcal{EL}$ w.r.t. static $\mathcal{EL}$ TBoxes.

The main novelty of this work is that we consider the relationship of contraction operations from belief change with *optimal* repairs (both in the classical and the general sense), i.e., repairs that are maximal subsets of the knowledge base to be repaired (classical case) or repairs that are entailed by the knowledge base to be repaired and preserve a maximal amount of consequences (general case). This notion of optimality usually does not play an important rôle in belief change (there is no optimality postulate), but under the assumption that the repair process should not lose consequences unnecessarily, it is important for ontology engineering. In [24, 25], classical repairs and gentle repairs are respectively set in relationship with contraction and pseudo-contraction operations, but optimal repairs are not considered. Work on revision and contraction for description logics [27] usually adapts the approaches from the belief change community to description logics as underlying logical formalism, but does not compare them with other ontology repair approaches, and in particular not with optimal repairs.

In recent work [29], the notion of a contraction has been introduced for $\mathcal{EL}$ concepts rather than knowledge bases. Basically, given two concepts C and D such that D is more general than C, a *concept contraction* of C w.r.t. D is a new concept C' such that C' is more general than C, but D is not more general than C'. The authors introduce concept contraction operations inspired by partial meet contractions, and characterize them using appropriate postulates. We will see that there is a close connection between such concept contractions and optimal repairs of the knowledge base $\{C(a)\}$ w.r.t. the unwanted consequence $D(a)$.

The next section introduces the general notion of a logical consequence operator, and then instantiates it with entailment from $\mathcal{EL}$ ABoxes w.r.t. an $\mathcal{EL}$ TBox. The definitions of contractions and repairs in the subsequent sections will be formulated in the general setting, with the concrete instance providing us with (counter)-examples. In Section 3, we first review relevant notions from belief change. In particular, we introduce partial meet and kernel contractions, and recall the postulates they satisfy. We then show that certain partial meet and kernel contractions always yield optimal classical repairs. Conversely, we note that a contraction operation that always returns an optimal classical repair (in case there is any repair) satisfies three of the four postulates characterizing partial meet contractions, but not the fourth (called *uniformity*). In Section 4, we introduce pseudo-contractions and in particular the "pseudo-versions" of partial meet and kernel contraction [30, 25]. Roughly speaking, we show that there always exists a partial meet pseudo-contraction that produces optimal repairs whenever such repairs exist, and optimal classical repairs otherwise. In general, however, partial meet pseudo-contractions need not yield optimal repairs (even if they exist) unless an additional property is satisfied. In Section 5, we recall the definitions and results for concept contractions of [29]. We relate concept contractions to optimal repairs of very simple knowledge bases, and show that this can be used to improve on the results obtained in [29] for this notion. Note that this section was not included in the earlier work [3] this work is based on. Section 6 summarizes the results achieved in this work and gives some hints regarding interesting future work.

2. PRELIMINARIES

Following [24], we assume that a logic is given by its language $\mathfrak{L}$, i.e., the set of sentences one can build in it, and its consequence operator $\mathsf{Cn} : 2^{\mathfrak{L}} \to 2^{\mathfrak{L}}$, which maps each set of sentences $\mathcal{X}$ to the set of its consequences $\mathsf{Cn}(\mathcal{X})$. Usually, $\mathfrak{L}$ will consist of certain first-order sentences, such as sentences expressed in some description logic, and Cn is first-order consequence restricted to $\mathfrak{L}$. Given sets of sentences $\mathcal{X}, \mathcal{Y} \subseteq \mathfrak{L}$ (a sentence $\alpha \in \mathfrak{L}$), we write $\mathcal{X} \models \mathcal{Y}$ ($\mathcal{X} \models \alpha$) to indicate that $\mathcal{Y} \subseteq \mathsf{Cn}(\mathcal{X})$ ($\alpha \in \mathsf{Cn}(\mathcal{X})$). In general, we only assume that Cn satisfies the following properties:

- $\mathcal{X} \subseteq \mathsf{Cn}(\mathcal{X})$ (inclusion),
- $\mathcal{X} \subseteq \mathcal{Y}$ implies $\mathsf{Cn}(\mathcal{X}) \subseteq \mathsf{Cn}(\mathcal{Y})$ (monotonicity),
- $\mathsf{Cn}(\mathsf{Cn}(\mathcal{X})) = \mathsf{Cn}(\mathcal{X})$ (idempotency),
- $\alpha \in \mathsf{Cn}(\mathcal{X})$ implies that there is a *finite* set $\mathcal{X}' \subseteq \mathcal{X}$ such that $\alpha \in \mathsf{Cn}(\mathcal{X}')$ (compactness).

These four properties are satisfied by first-order consequence, and thus also for most description logics.

As a concrete example, we consider ABoxes of the description logic $\mathcal{EL}$ as (finite) sets of sentences and consequence w.r.t. an $\mathcal{EL}$ TBox as the consequence operator. Our introduction of $\mathcal{EL}$ concepts, TBoxes, and ABoxes follows the presentation in [6].

The name space available for defining $\mathcal{EL}$ concepts and ABox assertions is given by a *signature* Σ, which is the disjoint union of sets Σ_{I}, Σ_{C}, and Σ_{R} of *individual names*, *concept* *names*, and *role names*. Starting with concept names and the top concept $\top$, $\mathcal{EL}$ *concepts* are defined inductively: if C, D are $\mathcal{EL}$ concepts and r is a role name, then $C \sqcap D$ (conjunction) and $\exists r. C$ (existential restriction) are also $\mathcal{EL}$ concepts. An $\mathcal{EL}$ *general concept inclusion (GCI)* is of the form $C \sqsubseteq D$, an $\mathcal{EL}$ *concept assertion* is of the form $C(a)$, and a *role assertion* is of the form $r(a, b)$, where C, D are $\mathcal{EL}$ concepts, $r \in \Sigma_{\mathsf{R}}$, and $a, b \in \Sigma_{\mathsf{I}}$. An $\mathcal{EL}$ *assertion* is a concept or a role assertion. An $\mathcal{EL}$ *TBox* is a finite set of $\mathcal{EL}$ GCIs and an $\mathcal{EL}$ *ABox* is a finite set of $\mathcal{EL}$ concept assertions and role assertions. Since, in this paper, we consider only one description logic, we sometimes omit the prefix $\mathcal{EL}$, and write assertion, ABox, etc. in place of $\mathcal{EL}$ assertion, $\mathcal{EL}$ ABox, etc.

The semantics of the syntactic entities introduced above can either be defined directly using interpretations or by a translation into first-order logic (FO) [4]. To make the connection to FO clearer, we choose here the latter approach. In the translation, the elements of Σ_{I}, Σ_{C}, and Σ_{R} are respectively viewed as constant symbols, unary predicate symbols, and binary predicate symbols. $\mathcal{EL}$ concepts C are inductively translated into FO formulas $\phi_C(x)$ with one free variable x:

- concept A for $A \in \Sigma_{\mathsf{C}}$ is translated into $A(x)$ and $\top$ into $A(x) \vee \neg A(x)$ for an arbitrary $A \in \Sigma_{\mathsf{C}}$;
- if C, D are translated into $\phi_C(x), \phi_D(x)$, then $C \sqcap D$ is translated into $\phi_C(x) \wedge \phi_D(x)$. The concept $\exists r. C$ is translated into $\exists y. (r(x, y) \wedge \phi_C(y))$, where $\phi_C(y)$ is obtained from $\phi_C(x)$ by replacing the free variable x by a variable y not occurring in $\phi_C(x)$.

GCIs $C \sqsubseteq D$ yield sentences $\phi_{C \sqsubseteq D} := \forall x. (\phi_C(x) \rightarrow \phi_D(x))$ and TBoxes $\mathcal{T}$ sets of sentences $\Phi_{\mathcal{T}} := \{\phi_{C \sqsubseteq D} \mid C \sqsubseteq D \in \mathcal{T}\}$. Concept assertions $C(a)$ are translated into $\phi_{C(a)} := \phi_C(a)$, role assertions $r(a, b)$ stay the same, i.e., $\phi_{r(a,b)} := r(a, b)$, and ABoxes $\mathcal{A}$ are translated into sets of sentences $\Phi_{\mathcal{A}} := \{\phi_\alpha \mid \alpha \in \mathcal{A}\}$.

The concept C is *subsumed* by the concept D w.r.t. the TBox $\mathcal{T}$ (written $C \sqsubseteq^{\mathcal{T}} D$) if $\phi_{C \sqsubseteq D}$ is a consequence of $\Phi_{\mathcal{T}}$ according to the semantics of FO. We write $C \equiv^{\mathcal{T}} D$ and say that C is *equivalent* to D w.r.t. $\mathcal{T}$ if they subsume each other w.r.t. $\mathcal{T}$. The concept C is *strictly subsumed* by the concept D w.r.t. the TBox $\mathcal{T}$ (written $C \sqsubset^{\mathcal{T}} D$) if $C \sqsubseteq^{\mathcal{T}} D$ and $C \not\equiv^{\mathcal{T}} D$.

The assertion α is a *consequence* of the set of assertions $\mathcal{A}$ w.r.t. the TBox $\mathcal{T}$ (written $\mathcal{A} \models^{\mathcal{T}} \alpha$) if ϕ_α is a consequence of the set of sentences $\Phi_{\mathcal{A}} \cup \Phi_{\mathcal{T}}$ according to the semantics of FO. This yields the consequence operator $\mathsf{Cn}_{\mathcal{T}}$, which takes as input a set of assertions $\mathcal{A}$, is parameterized with an $\mathcal{EL}$ TBox $\mathcal{T}$, and yields the following set of assertions as consequences:

$$\mathsf{Cn}_{\mathcal{T}}(\mathcal{A}) = \{\alpha \mid \mathcal{A} \models^{\mathcal{T}} \alpha \text{ where } \alpha \text{ is an } \mathcal{EL} \text{ assertion}\}.$$

Since its semantics is based on first-order consequence, $\mathsf{Cn}_{\mathcal{T}}$ clearly satisfies inclusion, monotonicity, idempotency, and compactness.

As an example, consider a situation where our rational agent believes that Ben has a parent called Jerry, who is both rich and famous. The agent also believes that people that have a rich and famous parent are arrogant. The former belief is represented in the ABox

$$\mathcal{A} := \{\, \mathit{has\_parent}(\mathit{BEN}, \mathit{JERRY}), \mathit{Famous}(\mathit{JERRY}), \mathit{Rich}(\mathit{JERRY}) \,\}$$

whereas the latter is expressed in the TBox

$$\mathcal{T} := \{\, \exists \mathit{has\_parent}. (\mathit{Famous} \sqcap \mathit{Rich}) \sqsubseteq \mathit{Arrogant} \,\}.$$

Clearly, we have $\mathit{Arrogant}(\mathit{BEN}) \in \mathsf{Cn}_{\mathcal{T}}(\mathcal{A})$. Now assume that the agent actually meets Ben and notices that he is not arrogant. Since the agent insists on sticking with the prejudice that children of rich and famous people are arrogant, the unwanted consequence $\mathit{Arrogant}(\mathit{BEN})$ can only be removed by modifying the ABox. In the classical repair approach, this can be achieved by removing one of its three assertions from $\mathcal{A}$. Let us assume that the agent decides to remove $\mathit{Famous}(\mathit{JERRY})$. This removes the unwanted consequence $\mathit{Arrogant}(\mathit{BEN})$, but also the consequence $\exists \mathit{has\_parent}. \mathit{Famous}(\mathit{BEN})$.

Removing $\mathit{Famous}(\mathit{JERRY})$ from $\mathcal{A}$, but adding the assertion $\exists \mathit{has\_parent}. \mathit{Famous}(\mathit{BEN})$ to the ABox yields a repair that retains more consequences than the classical repair. This improved repair corresponds to the agent's new belief that Jerry is only rich, and that Ben has another famous parent, whose name is not known to the agent.

3. CLASSICAL REPAIRS AND CONTRACTIONS

The classical notions of contraction and repair resort to subsets of the given knowledge base to remove an unwanted consequence. Following [30, 24, 25], we first define contractions and recall two approaches for constructing them. Then, we describe their connection to classical repairs.

3.1 Contractions in Belief Change

Let $\mathfrak{L}$ be a logical language and Cn a monotone, idempotent, and compact consequence operator satisfying inclusion. A *belief base* is an arbitrary subset of $\mathfrak{L}$. Contractions get rid of unwanted consequences of a belief base by removing some of its sentences. More formally, a *contraction operation* ctr accepts a belief base $\mathcal{B} \subseteq \mathfrak{L}$ and a sentence $\alpha \in \mathfrak{L}$ as input, and produces as output a belief base $\mathsf{ctr}(\mathcal{B}, \alpha)$ that satisfies the following two postulates:

- $\mathsf{ctr}(\mathcal{B}, \alpha) \subseteq \mathcal{B}$ (inclusion),
- if $\alpha \notin \mathsf{Cn}(\emptyset)$, then $\alpha \notin \mathsf{Cn}(\mathsf{ctr}(\mathcal{B}, \alpha))$ (success).

In the belief change literature, reasonable contraction operations are usually assumed to satisfy additional postulates. This is the case for contractions obtained by applying one of the following two prominent approaches for constructing contraction operations: partial meet contraction [1, 18] and kernel contraction [19]. To define the former, we must introduce remainders, remainder sets, and selection functions. Let $\mathcal{B}$ be a belief base and α a sentence.

- A *remainder* of $\mathcal{B}$ with respect to α is a maximal subset $\mathcal{X}$ of $\mathcal{B}$ such that $\alpha \notin \mathsf{Cn}(\mathcal{X})$. We denote the set of all remainders of $\mathcal{B}$ with respect to α as $\mathsf{rem}(\mathcal{B}, \alpha)$.

- A *selection function* γ for $\mathcal{B}$ takes such sets of remainders as input and satisfies the following properties for each $\alpha \in \mathfrak{L}$:
  - $\emptyset \neq \gamma(\mathsf{rem}(\mathcal{B},\alpha)) \subseteq \mathsf{rem}(\mathcal{B},\alpha)$ if $\mathsf{rem}(\mathcal{B},\alpha) \neq \emptyset$,
  - $\gamma(\mathsf{rem}(\mathcal{B},\alpha)) = \{\mathcal{B}\}$ if $\mathsf{rem}(\mathcal{B},\alpha) = \emptyset$.

Note that the value returned by the selection function does not depend of α itself, but on the set $\mathsf{rem}(\mathcal{B},\alpha)$. In case this set is non-empty, this value is a non-empty subset of $\mathsf{rem}(\mathcal{B},\alpha)$. Otherwise, the set consisting of $\mathcal{B}$ is returned. This second case occurs iff $\alpha \in \mathsf{Cn}(\emptyset)$. Each selection function γ induces a *partial meet contraction* operation ctr_γ as follows:

$$\mathsf{ctr}_\gamma(\mathcal{B},\alpha) := \bigcap \gamma(\mathsf{rem}(\mathcal{B},\alpha)).$$

As shown by Hansson in [18], the operation ctr_γ satisfies *inclusion* and *success*, and thus is a contraction operation, and additionally the following postulates:

- if $\beta \in \mathcal{B}\backslash\mathsf{ctr}(\mathcal{B},\alpha)$, then there is $\mathcal{B}'$ such that $\mathsf{ctr}(\mathcal{B},\alpha) \subseteq \mathcal{B}' \subseteq \mathcal{B}$, $\alpha \notin \mathsf{Cn}(\mathcal{B}')$, and $\alpha \in \mathsf{Cn}(\mathcal{B}' \cup \{\beta\})$ (relevance),
- if $\alpha \in \mathsf{Cn}(\mathcal{B}')$ iff $\beta \in \mathsf{Cn}(\mathcal{B}')$ holds for all $\mathcal{B}' \subseteq \mathcal{B}$, then $\mathsf{ctr}(\mathcal{B},\alpha) = \mathsf{ctr}(\mathcal{B},\beta)$ (uniformity).

Hansson [18] also shows that any contraction operation that satisfies the postulates *inclusion*, *success*, *relevance*, and *uniformity* can be obtained as a partial meet contraction. In [19] he introduces another construction for obtaining contraction operations, which is based on the notions of kernels and incision functions.

- The *kernel* $\mathsf{ker}(\mathcal{B},\alpha)$ of $\mathcal{B}$ with respect to α consists of the minimal subsets $\mathcal{X}$ of $\mathcal{B}$ satisfying $\alpha \in \mathsf{Cn}(\mathcal{X})$.
- An *incision function* σ for $\mathcal{B}$ takes such kernel sets as input and satisfies the following properties for each $\alpha \in \mathcal{B}$:
  - $\sigma(\mathsf{ker}(\mathcal{B},\alpha)) \subseteq \bigcup \mathsf{ker}(\mathcal{B},\alpha)$,
  - If $\mathcal{X}$ is a non-empty element of $\mathsf{ker}(\mathcal{B},\alpha)$, then $\mathcal{X} \cap \sigma(\mathsf{ker}(\mathcal{B},\alpha)) \neq \emptyset$.

Like selection functions, incision functions depend only on the kernel set $\mathsf{ker}(\mathcal{B},\alpha)$, and not on the sentence α itself. It is easy to see that $\emptyset \in \mathsf{ker}(\mathcal{B},\alpha)$ iff $\mathsf{ker}(\mathcal{B},\alpha) = \{\emptyset\}$ iff $\alpha \in \mathsf{Cn}(\emptyset)$. Each incision function σ induces a *kernel contraction* operation ctr_σ as follows:

$$\mathsf{ctr}_\sigma(\mathcal{B},\alpha) := \mathcal{B} \setminus \sigma(\mathsf{ker}(\mathcal{B},\alpha)).$$

As shown by Hansson in [19], the operation ctr_σ satisfies *inclusion*, *success*, and *uniformity*, but *relevance* needs to be replaced by the following weaker postulate:

- if $\beta \in \mathcal{B} \setminus \mathsf{ctr}(\mathcal{B},\alpha)$, then there is $\mathcal{B}' \subseteq \mathcal{B}$ such that $\alpha \notin \mathsf{Cn}(\mathcal{B}')$ and $\alpha \in \mathsf{Cn}(\mathcal{B}' \cup \{\beta\})$ (core-retainment).

Any contraction operation that satisfies the postulates *inclusion*, *success*, *core-retainment*, and *uniformity* can be obtained as a kernel contraction [19].

3.2 Classical Repairs in Ontology Engineering

Knowledge bases in ontology engineering are usually assumed to be finite. Thus, given a logical language $\mathfrak{L}$ and a monotone, idempotent, and compact consequence operator Cn satisfying inclusion, a *knowledge base* is a *finite* subset of $\mathfrak{L}$.

A classical repair is then just a contraction, i.e., given a knowledge base $\mathcal{B} \subseteq \mathfrak{L}$ and a sentence $\alpha \in \mathfrak{L}$, a *classical repair* of $\mathcal{B}$ with respect to α is (by definition) a subset $\mathcal{X}$ of $\mathcal{B}$ that satisfies $\alpha \notin \mathsf{Cn}(\mathcal{X})$ [8]. Thus, if we consider an operation $\mathsf{ctr}_{\mathsf{rep}}$ that, on input $\mathcal{B}$ and α, returns a classical repair of $\mathcal{B}$ with respect to α if $\alpha \notin \mathsf{Cn}(\emptyset)$, and $\mathcal{B}$ otherwise, then $\mathsf{ctr}_{\mathsf{rep}}$ satisfies *inclusion* and *success*, and thus is a contraction operation (see Proposition 3 in [24]).

In ontology engineering, one usually wants to remove a minimal amount of information to eliminate an unwanted consequence. Thus, one is interested in computing optimal classical repairs. Given a knowledge base $\mathcal{B} \subseteq \mathfrak{L}$ and a sentence $\alpha \in \mathfrak{L}$, an *optimal classical repair* of $\mathcal{B}$ with respect to α is a *maximal* subset $\mathcal{X}$ of $\mathcal{B}$ satisfying $\alpha \notin \mathsf{Cn}(\mathcal{X})$. Obviously, the notions *optimal classical repair* and *remainder* coincide, which yields the following proposition.

PROPOSITION 1. *Let $\mathcal{B}$ be a knowledge base, α a sentence, and γ a selection function for $\mathcal{B}$ such that $|\gamma(\mathsf{rem}(\mathcal{B},\alpha))| = 1$ for all $\alpha \in \mathfrak{L}$. Then $\mathsf{ctr}_\gamma(\mathcal{B},\alpha)$ is an optimal classical repair of $\mathcal{B}$ with respect to α for all sentences α satisfying $\alpha \notin \mathsf{Cn}(\emptyset)$, and $\mathsf{ctr}_\gamma(\mathcal{B},\alpha) = \mathcal{B}$ if $\alpha \in \mathsf{Cn}(\emptyset)$.*

In [1, 24], a partial meet contraction operation defined using a selection function γ satisfying $|\gamma(\mathsf{rem}(\mathcal{B},\alpha))| = 1$ for all sentences α is called a *maxichoice* contraction operation. Thus, one can rephrase the statement of Proposition 1 as follows.

COROLLARY 2. *If ctr is a maxichoice contraction operation and $\mathcal{B}$ has a classical repair with respect to α, then $\mathsf{ctr}(\mathcal{B},\alpha)$ is an optimal classical repair of $\mathcal{B}$ with respect to α.*

Optimal classical repairs can also be obtained as kernel contractions. In fact, in ontology engineering, optimal classical repairs are often constructed using justifications and Reiter's hitting set duality [28]. Before we can describe this approach, we must introduce the relevant notions. Let $\mathcal{B}$ be a knowledge base and α a sentence.

- A *justification* of α in $\mathcal{B}$ is a minimal subset $\mathcal{X}$ of $\mathcal{B}$ such that $\alpha \in \mathsf{Cn}(\mathcal{X})$. We denote the set of all justifications of α in $\mathcal{B}$ as $\mathsf{jus}(\mathcal{B},\alpha)$. Note that $\mathsf{jus}(\mathcal{B},\alpha) = \emptyset$ if $\alpha \notin \mathsf{Cn}(\mathcal{B})$, and $\mathsf{jus}(\mathcal{B},\alpha) = \{\emptyset\}$ if $\alpha \in \mathsf{Cn}(\emptyset)$.
- Given a collection $\{\mathcal{X}_1,\ldots,\mathcal{X}_k\}$ of subsets $\mathcal{X}_i$ of $\mathcal{B}$, a *hitting set* $\mathcal{H}$ of this collection is a subset of $\mathcal{X}_1 \cup \ldots \cup \mathcal{X}_k$ such that $\mathcal{H} \cap \mathcal{X}_i \neq \emptyset$ for all $i = 1,\ldots,k$. This hitting set is *minimal* if no other hitting set is strictly contained in it. Note: if the collection is empty (i.e., if $k = 0$), then $\emptyset$ is a minimal hitting set; if it contains the empty set (i.e., if $\mathcal{X}_i = \emptyset$ for some $i, 1 \leq i \leq k$), then it has no hitting set.

It is well-known [28, 8] that the optimal classical repairs of $\mathcal{B}$ with respect to α are exactly the sets $\mathcal{B} \setminus \mathcal{H}$ where $\mathcal{H}$ ranges over the minimal hitting sets of $\mathsf{jus}(\mathcal{B}, \alpha)$. Note that this characterization also works in the following borderline cases. If $\alpha \in \mathsf{Cn}(\emptyset)$, then there is no optimal classical repair, and neither is there a hitting set of $\mathsf{jus}(\mathcal{B}, \alpha) = \{\emptyset\}$. If $\alpha \notin \mathsf{Cn}(\mathcal{B})$, then $\mathcal{B}$ is the only optimal classical repair, and $\mathsf{jus}(\mathcal{B}, \alpha) = \emptyset$ has $\emptyset$ as its only minimal hitting set.

Obviously, the set of all justifications of α in $\mathcal{B}$ coincides with $\mathsf{ker}(\mathcal{B}, \alpha)$, i.e., $\mathsf{jus}(\mathcal{B}, \alpha) = \mathsf{ker}(\mathcal{B}, \alpha)$. In addition, if $\alpha \notin \mathsf{Cn}(\emptyset)$, then $\sigma(\mathsf{ker}(\mathcal{B}, \alpha))$ is a hitting set of $\mathsf{jus}(\mathcal{B}, \alpha) = \mathsf{ker}(\mathcal{B}, \alpha)$ for every incision function σ. We call an incision function *minimal* if $\sigma(\mathsf{ker}(\mathcal{B}, \alpha))$ is a minimal hitting set of $\mathsf{ker}(\mathcal{B}, \alpha)$ for all α with $\alpha \notin \mathsf{Cn}(\emptyset)$, and $\sigma(\mathsf{ker}(\mathcal{B}, \alpha)) = \emptyset$ if $\alpha \in \mathsf{Cn}(\emptyset)$.

PROPOSITION 3. *Let $\mathcal{B}$ be a knowledge base, α a sentence, and σ a minimal incision function for $\mathcal{B}$. Then $\mathsf{ctr}_\sigma(\mathcal{B}, \alpha)$ is an optimal classical repair of $\mathcal{B}$ with respect to α for all sentences α satisfying $\alpha \notin \mathsf{Cn}(\emptyset)$, and $\mathsf{ctr}_\sigma(\mathcal{B}, \alpha) = \mathcal{B}$ if $\alpha \in \mathsf{Cn}(\emptyset)$.*

Using Reiter's hitting set duality [28], it is easy to see that every maxichoice partial meet contraction can be obtained as a kernel contraction induced by a minimal incision function, and vice versa (see [12] for details).

Now, consider a special case $\mathsf{ctr}_{\mathsf{orep}}$ of the contraction operation $\mathsf{ctr}_{\mathsf{rep}}$ introduced above, where we require that $\mathsf{ctr}_{\mathsf{orep}}(\mathcal{B}, \alpha)$ is an *optimal* classical repair of $\mathcal{B}$ with respect to α if $\alpha \notin \mathsf{Cn}(\emptyset)$.

PROPOSITION 4. *The operation $\mathsf{ctr}_{\mathsf{orep}}$ satisfies* inclusion*,* success*, and* relevance*, but it need not satisfy* uniformity*.*

PROOF. We already know that *inclusion* and *success* are satisfied even in the more general setting where an arbitrary repair, rather than an optimal one, is chosen. To show *relevance*, assume that $\beta \in \mathcal{B} \setminus \mathsf{ctr}_{\mathsf{orep}}(\mathcal{B}, \alpha)$. If we take $\mathcal{B}' := \mathsf{ctr}_{\mathsf{orep}}(\mathcal{B}, \alpha)$, then $\mathsf{ctr}_{\mathsf{orep}}(\mathcal{B}, \alpha) \subseteq \mathcal{B}' \subseteq \mathcal{B}$ is satisfied. Maximality of $\mathsf{ctr}_{\mathsf{orep}}(\mathcal{B}, \alpha)$ yields $\alpha \in \mathsf{Cn}(\mathcal{B}' \cup \{\beta\})$.

Without additional assumptions on how the optimal repairs are chosen, *uniformity* need not be satisfied. This is demonstrated by the example presented below. □

EXAMPLE 5. Consider the logical language that consists of $\mathcal{EL}$ assertions and the consequence operator $\mathsf{Cn}_\mathcal{T}$ for the $\mathcal{EL}$ TBox $\mathcal{T} := \{A \sqcap B \sqsubseteq C, A \sqcap B \sqsubseteq D\}$, and set $\mathcal{B} := \{A(a), B(a)\}$, $\alpha := C(a)$, and $\beta := D(a)$. Then $\alpha \in \mathsf{Cn}_\mathcal{T}(\mathcal{B}')$ iff $\beta \in \mathsf{Cn}_\mathcal{T}(\mathcal{B}')$ holds for all $\mathcal{B}' \subseteq \mathcal{B}$. In fact, for $\mathcal{B}' = \mathcal{B}$, both α and β belong to $\mathsf{Cn}_\mathcal{T}(\mathcal{B}')$, whereas for $\mathcal{B}' \subset \mathcal{B}$ neither α nor β belongs to $\mathsf{Cn}_\mathcal{T}(\mathcal{B}')$. However, our contraction operation $\mathsf{ctr}_{\mathsf{orep}}$ could choose the optimal classical repair $\{A(a)\}$ for α and $\{B(a)\}$ for β, thus violating *uniformity*.

The problem in this example is caused by the fact that α and β produce the same sets of optimal classical repairs. If in such a situation we insist that $\mathsf{ctr}_{\mathsf{orep}}$ chooses the same element of this set for both α and β, then $\mathsf{ctr}_{\mathsf{orep}}$ is actually a maxichoice partial meet contraction, and thus also satisfies *uniformity*.

4. OPTIMAL REPAIRS AND PSEUDO-CONTRACTIONS

The classical notions of contraction and repair have the disadvantage that they are syntax-dependent in the sense that a contraction (repair) can only use the sentences that are explicitly present in the belief (knowledge) base. This may lead to removal of more consequences than is necessary to get rid of the unwanted one. For example, consider the ABoxes $\mathcal{A} := \{(A \sqcap B)(a)\}$ and $\mathcal{B} := \{(A(a), B(a)\}$, and let the unwanted consequence be $\alpha := A(a)$. The two ABoxes $\mathcal{A}$, $\mathcal{B}$ are equivalent (i.e., $\mathsf{Cn}_\emptyset(\mathcal{A}) = \mathsf{Cn}_\emptyset(\mathcal{B})$). However, with respect to α, the ABox $\mathcal{A}$ has the empty ABox as only optimal classical repair for the consequence operator $\mathsf{Cn}_\emptyset$, whereas $\mathcal{B}$ has the optimal classical repair $\{B(a)\}$. Thus, the latter repair retains the consequence $B(a)$, whereas the former does not. Pseudo-contractions and optimal repairs try to overcome this problem.

4.1 Pseudo-Contractions in Belief Change

The problem of syntax-dependency is caused by the *inclusion* postulate. In the definition of pseudo-contractions, this postulate is replaced by *logical inclusion* [15, 17]:

- $\mathsf{Cn}(\mathsf{ctr}(\mathcal{B}, \alpha)) \subseteq \mathsf{Cn}(\mathcal{B})$ (logical inclusion).

The operation $\mathsf{ctr} : 2^\mathfrak{L} \times \mathfrak{L} \to 2^\mathfrak{L}$ is a *pseudo-contraction* operation if it satisfies *success* and *logical inclusion*.

To construct pseudo-contractions that retain more consequences than contractions, one can first add some of the logical consequences of $\mathcal{B}$ to the given belief base $\mathcal{B}$, and then apply the partial meet or the kernel contraction approach to the resulting extended belief base [30, 24, 25]. In the cited literature, both one-place and two-place extension functions Cn^* are considered, where the former add consequences independently of the unwanted sentence α, whereas the latter also take α into account. Here, we consider the two-place setting since it makes it easier to obtain a connection with optimal repairs. A *two-place consequence operator* is a function $\mathsf{Cn}^* : 2^\mathcal{L} \times \mathcal{L} \to 2^\mathcal{L}$. We call such an operator a *two-place extension function* with respect to Cn if it satisfies $\mathcal{B} \subseteq \mathsf{Cn}^*(\mathcal{B}, \alpha) \subseteq \mathsf{Cn}(\mathcal{B})$ for all belief bases $\mathcal{B}$ and sentences α. This operator is further called *finite* if $\mathsf{Cn}^*(\mathcal{B}, \alpha)$ is finite whenever its first argument $\mathcal{B}$ is finite. In case the value returned by Cn^* does not depend on the second argument, we write $\mathsf{Cn}^*(\mathcal{B})$ in place of $\mathsf{Cn}^*(\mathcal{B}, \alpha)$ and call Cn^* a *one-place extension function* with respect to Cn.

EXAMPLE 6. In the $\mathcal{EL}$ ABox setting, one can define a one-place extension function with respect to $\mathsf{Cn}_\mathcal{T}$ by breaking conjunctions in concept assertions into their conjuncts, i.e., if $C(a) \in \mathcal{B}$ and $C = C_1 \sqcap \ldots \sqcap C_n$ where the C_i are existential restrictions or concept names, then $C_1(a), \ldots, C_n(a)$ are added to $\mathsf{Cn}^*(\mathcal{B})$ (see Example 6 in [25]). Clearly, this yields a finite extension function with respect to $\mathsf{Cn}_\mathcal{T}$. In our introductory example, for $\mathcal{A} = \{(A \sqcap B)(a)\}$, we obtain $\mathsf{Cn}^*_\emptyset(\mathcal{A}) = \mathcal{A} \cup \{(A(a), B(a)\}$.
Another possibility is to add assertions entailed by the TBox. To keep the extension function finite, we can, e.g., restrict this to concept assertions for concept names: for every concept name $A \in \Sigma_\mathsf{C}$ and every individual name a occurring

in $\mathcal{B}$, add $A(a)$ to $\mathsf{Cn}^*(\mathcal{B})$ if $\mathcal{B} \models^{\mathcal{T}} A(a)$. This yields a finite extension function since it is easy to see that $A(a)$ can only be entailed if the concept name A occurs in $\mathcal{B}$ or $\mathcal{T}$.

The idea is now to apply the partial meet or the kernel contraction approach to $\mathsf{Cn}^*(\mathcal{B}, \alpha)$ rather than to $\mathcal{B}$. A Cn^* *partial meet pseudo-contraction* is thus obtained by considering remainders and selection functions of $\mathsf{Cn}^*(\mathcal{B}, \alpha)$. Given the set of remainders $\mathsf{rem}(\mathsf{Cn}^*(\mathcal{B}, \alpha), \alpha)$ and a selection function γ^* of $\mathsf{Cn}^*(\mathcal{B}, \alpha)$, the Cn^* partial meet pseudo-contraction induced by γ^* is then defined as

$$ctr^*_{\gamma^*}(\mathcal{B}, \alpha) := \bigcap \gamma^*(\mathsf{rem}(\mathsf{Cn}^*(\mathcal{B}, \alpha), \alpha)).$$

For the ABox $\mathcal{A} = \{(A \sqcap B)(a)\}$ of Example 6, the only remainder of $\mathsf{Cn}^*_\emptyset(\mathcal{A}) = \mathcal{A} \cup \{(A(a), B(a)\}$ with respect to $\alpha = A(a)$ is $\{B(a)\}$, and thus the selection function γ^* must choose this remainder. This shows that $\mathsf{ctr}^*_{\gamma^*}(\mathcal{A}, \alpha) = \{B(a)\}$.

Cn^* *kernel pseudo-contractions* are defined analogously, by using kernels and incision functions for $\mathsf{Cn}^*(\mathcal{B}, \alpha)$ rather than for $\mathcal{B}$. In the example, the kernel set of $\mathsf{Cn}^*_\emptyset(\mathcal{B}, \alpha)$ consists of the sets $\{(A \sqcap B)(a)\}$ and $\{A(a)\}$, and thus the only hitting set is $\{(A \sqcap B)(a), (A(a)\}$, which thus must be chosen by the incision function δ^*. This shows that the $\mathsf{Cn}^*_\emptyset$ kernel pseudo-contractions $\mathsf{ctr}^*_{\delta^*}(\mathcal{A}, \alpha)$ is in this case also $\{B(a)\}$.

Basically, these pseudo-contractions inherit the postulates satisfied by the underlying contraction operations, but they need to be formulated in a "starred" variant that takes the application of Cn^* into account, and they may depend also on properties of Cn^* (like monotonicity). More details regarding postulates can be found in [30, 24, 25]. Here, we only point out that, as an obvious consequence of the definition of *extension function* and the fact that kernel and partial meet contractions satisfy *inclusion* and *success*, the Cn^* kernel and partial meet pseudo-contractions introduced above satisfy *logical inclusion* and *success*, and thus are indeed pseudo-contractions.

4.2 Repairs and Optimal Repairs in Knowledge Engineering

Given a knowledge base $\mathcal{A}$ and a sentence α, a *repair* of $\mathcal{A}$ with respect to α is a knowledge base $\mathcal{B}$ that satisfies $\mathcal{B} \subseteq \mathsf{Cn}(\mathcal{A})$ and $\alpha \notin \mathsf{Cn}(\mathcal{B})$ [8]. Thus, like pseudo-contractions, repairs must satisfy *logical inclusion* and *success*. Since the repair must again be a knowledge base, a pseudo-contraction ctr only yields repairs if it additionally satisfies the following postulate:

- if $\mathcal{B}$ is finite, then $\mathsf{ctr}(\mathcal{B}, \alpha)$ is also finite (finiteness).

Contractions satisfy finiteness since they yield a subset of the input set $\mathcal{B}$. Since Cn^* partial meet or kernel pseudo-contractions yield contractions of $\mathsf{Cn}^*(\mathcal{B}, \alpha)$, their output is finite if $\mathsf{Cn}^*(\mathcal{B}, \alpha)$ is finite.

PROPOSITION 7. *If Cn^* is finite, $\mathcal{B}$ is a knowledge base, and α is a sentence, then $\mathsf{ctr}(\mathcal{B}, \alpha)$ is a repair whenever ctr is a Cn^* partial meet or kernel pseudo-contraction.*

To obtain repairs that preserve more consequences than classical repairs, an approach similar to the one described in the previous subsection is used, e.g., in [20, 11]. In these papers, a specific syntactic structural transformation is applied to the axioms in an ontology, which replaces them by sets of logically weaker axioms. The knowledge bases obtained by this approach are then repaired using the classical approach. There are also repair methods that directly apply weakening operations to axioms to construct a repair, such as the ones described in [23, 32, 8]. The connection between such "gentle repairs" and pseudo-contractions has been investigated in [24, 25].

Here, we concentrate on optimal repairs instead. Given a knowledge base $\mathcal{A}$ and a sentence α, the repair $\mathcal{B}$ of $\mathcal{A}$ with respect to α is *optimal* if there is no repair $\mathcal{C}$ of $\mathcal{A}$ with respect to α such that $\mathcal{C} \models \mathcal{B}$ and $\mathcal{B} \not\models \mathcal{C}$ [8]. As shown in [8], optimal repairs need not exists even if there are repairs.

EXAMPLE 8 ([8]). Consider the logical language that consists of $\mathcal{EL}$ assertions and the consequence operator $\mathsf{Cn}_{\mathcal{T}}$ for the $\mathcal{EL}$ TBox $\mathcal{T} := \{A \sqsubseteq \exists r.A, \exists r.A \sqsubseteq A\}$, and set $\mathcal{A} := \{A(a)\}$ and $\alpha := A(a)$. The empty ABox is clearly a repair in this case. However, as shown in the proof of Proposition 2 in [8], $\mathcal{A}$ does not have an optimal repair. Intuitively, the reason for this is that any ABox of the form $\mathcal{A}_n := \{(\exists r.)^n \top(a)\}$ for $n \geq 1$ is a repair, but any fixed repair can entail only finitely many of them. Thus, if $\mathcal{B}$ is a repair, then there is an n such that $\mathcal{B} \not\models^{\mathcal{T}} \mathcal{A}_n$. But then $\mathcal{B} \cup \mathcal{A}_n$ is a repair that entails $\mathcal{B}$, but is not entailed by $\mathcal{B}$, which shows that $\mathcal{B}$ cannot be optimal.

Moreover, even if optimal repairs exist, they need not cover all repairs in a sense to be made more precise below. First, note that optimal classical repairs cover all classical repairs in the sense that every classical repair is contained in an optimal classical repair. For general repairs, the notion of containment needs to be replaced by entailment, i.e., containment of the consequence sets. We say that the set of all optimal repairs of $\mathcal{A}$ with respect to α *covers* all repairs of $\mathcal{A}$ with respect to α if, for every repair $\mathcal{B}$ of $\mathcal{A}$ with respect to α, there is an optimal repair $\mathcal{C}$ of $\mathcal{A}$ with respect to α such that $\mathcal{C} \models \mathcal{B}$.

EXAMPLE 9 ([6]). Consider the TBox $\mathcal{T}$, ABox $\mathcal{A}$, and sentence α, where $\mathcal{T} := \{B \sqsubseteq \exists r.B, \exists r.B \sqsubseteq B\}$, $\mathcal{A} := \{A(a), r(a, b), B(b)\}$, and $\alpha := (A \sqcap \exists r.B)(a)$. As shown in [6] (Example 12), for the consequence operator $\mathsf{Cn}_{\mathcal{T}}$, the ABox $\mathcal{C} := \{r(a, b), B(b)\}$ is the only optimal repair of $\mathcal{A}$ with respect to α. However, the ABox $\mathcal{B} := \{A(a), r(a, b), (\exists r.\exists r.\top)(b)\}$ is also a repair of $\mathcal{A}$ with respect to α, but it is not entailed by $\mathcal{C}$. Thus, in this example, the set of optimal repairs does not cover all repairs.

In the remainder of this section, we investigate the connection between optimal repairs and partial meet and kernel pseudo-contractions. For this, we first need to define an appropriate consequence operator Cn^*. Let $\mathcal{A}$ be a knowledge base and α a sentence. We define $\mathsf{Orep}(\mathcal{A}, \alpha)$ to consists of the optimal repairs of $\mathcal{A}$ with respect to α,[1] and set

$$\mathsf{Cn}^*(\mathcal{A}, \alpha) := \mathcal{A} \cup \bigcup \mathsf{Orep}(\mathcal{A}, \alpha).$$

[1] More precisely, we assume that $\mathsf{Orep}(\mathcal{A}, \alpha)$ contains one

This operator is a two-place extension function w.r.t. Cn since $\mathcal{A} \subseteq \mathsf{Cn}^*(\mathcal{A}, \alpha)$ by definition and $\mathsf{Cn}^*(\mathcal{A}, \alpha) \subseteq \mathsf{Cn}(\mathcal{A})$ holds because every repair of $\mathcal{A}$ is entailed by $\mathcal{A}$. This extension function is finite iff $\mathsf{Orep}(\mathcal{A}, \alpha)$ is finite for all knowledge bases $\mathcal{A}$ and sentences α. This condition is satisfied in our ABox setting.

PROPOSITION 10. *Let $\mathcal{T}$ be an $\mathcal{EL}$ TBox and $\mathsf{Cn}_\mathcal{T}$ the induced consequence operator on $\mathcal{EL}$ ABoxes. Then $\mathsf{Cn}^*_\mathcal{T}$ is a finite extension function that can effectively be computed.*

PROOF. It remains to show that $\mathsf{Cn}^*_\mathcal{T}$ is finite and computable. This is an easy consequence of the results proved in [6]. In fact, it is shown there that the optimal ABox repairs of $\mathcal{A}$ with respect to α can be computed by first computing the optimal quantified ABox (qABox) repairs of $\mathcal{A}$ with respect to α for IRQ-entailment. This set is finite and can effectively be computed. The optimal ABox repairs are obtained from this set by computing, for each qABox in this set, its optimal ABox approximation, if it exists. Existence of this approximation is decidable, and if it exists, then the approximation can be computed. □

The next lemma yields a connection between optimal repairs and the notion of a remainder.

LEMMA 11. *Let $\mathcal{A}$ be a knowledge base and α a sentence. If $\mathcal{B} \in \mathsf{Orep}(\mathcal{A}, \alpha)$, then $\mathcal{B}$ is equivalent to a remainder of $\mathsf{Cn}^*(\mathcal{A}, \alpha)$ with respect to α.*

PROOF. We must show that $\mathcal{B} \in \mathsf{Orep}(\mathcal{A}, \alpha)$ is equivalent to a maximal subset of $\mathsf{Cn}^*(\mathcal{A}, \alpha)$ that does not have the consequence α. Since it is a repair with respect to α, it does not have the consequence α. Assume that $\mathcal{B}$ is not maximal, i.e., there is $\mathcal{B} \subset \mathcal{B}' \subseteq \mathsf{Cn}^*(\mathcal{A}, \alpha)$ such that $\alpha \notin \mathsf{Cn}(\mathcal{B}')$. We can assume without loss of generality that $\mathcal{B}'$ is a remainder.[2] If $\mathcal{B}'$ is equivalent to $\mathcal{B}$, then we are done. Otherwise, we obtain a contradiction to our assumption that $\mathcal{B}$ is an optimal repair. □

The following result is an easy consequence of this lemma.

THEOREM 12. *Let $\mathcal{A}$ be a knowledge base and α a sentence. There exists a Cn^* partial meet pseudo-contraction $\mathsf{ctr}^*_{\gamma^*}$ such that $\mathsf{ctr}^*_{\gamma^*}(\mathcal{A}, \alpha)$ is an optimal repair of $\mathcal{A}$ w.r.t. α if $\mathsf{Orep}(\mathcal{A}, \alpha) \neq \emptyset$, and an optimal classical repair of $\mathcal{A}$ w.r.t. α if $\mathsf{Orep}(\mathcal{A}, \alpha) = \emptyset$ and $\alpha \notin \mathsf{Cn}(\emptyset)$.*

PROOF. Define γ^* such that it chooses an element of the set $\mathsf{Orep}(\mathcal{A}, \alpha)$ if it is non-empty, and an arbitrary remainder of $\mathsf{Cn}^*(\mathcal{A}, \alpha)$ otherwise. By Lemma 11, this indeed yields a selection function for $\mathsf{Cn}^*(\mathcal{A}, \alpha)$. In case $\mathsf{Orep}(\mathcal{A}, \alpha) = \emptyset$, we know that $\mathsf{Cn}^*(\mathcal{A}, \alpha) = \mathcal{A}$, and thus a remainder is an optimal classical repair in this case, unless there is no repair. □

representative of every equivalence class of optimal repairs, where two knowledge bases are equivalent if they entail each other.

[2]If $\mathsf{Cn}^*(\mathcal{A}, \alpha)$ is finite, then this is trivial. Otherwise, one needs to use transfinite induction and the fact that Cn is compact.

In general, remainders of $\mathsf{Cn}^*(\mathcal{A}, \alpha)$ need not be optimal repairs even if $\mathsf{Orep}(\mathcal{A}, \alpha) \neq \emptyset$.

EXAMPLE 13. Consider the ABox $\mathcal{A}$, the TBox $\mathcal{T}$, and the sentence α of Example 9. Since in this case the only optimal repair is a subset of $\mathcal{A}$, we have $\mathsf{Cn}^*_\mathcal{T}(\mathcal{A}, \alpha) = \mathcal{A}$. The ABox $\mathcal{B}' := \{A(a), r(a, b)\}$ is a remainder of $\mathsf{Cn}^*_\mathcal{T}(\mathcal{A}, \alpha)$, but it is not optimal since $\mathcal{B} = \{A(a), r(a, b), (\exists r.\exists r.\top)(b)\}$ is a repair that strictly entails $\mathcal{B}'$.

This problem cannot occur if $\mathsf{Orep}(\mathcal{A}, \alpha)$ covers all repairs.

LEMMA 14. *Let $\mathcal{A}$ be a knowledge base and α a sentence such that $\mathsf{Orep}(\mathcal{A}, \alpha)$ covers all repairs of $\mathcal{A}$ w.r.t. α. If $\mathcal{B}$ is a remainder of $\mathsf{Cn}^*(\mathcal{A}, \alpha)$ w.r.t. α, then $\mathcal{B}$ is an optimal repair of $\mathcal{A}$ w.r.t. α.*

PROOF. First, note that $\mathcal{B}$ is entailed by $\mathcal{A}$ since $\mathcal{B} \subseteq \mathsf{Cn}^*(\mathcal{A}, \alpha) \subseteq \mathsf{Cn}(\mathcal{A})$. In addition, $\alpha \notin \mathsf{Cn}(\mathcal{B})$ holds by the definition of a remainder. Thus, $\mathcal{B}$ is a repair of $\mathcal{A}$ with respect to α.

Assume that $\mathcal{B}$ is not optimal. Then there is a repair $\mathcal{B}'$ of $\mathcal{A}$ with respect to α that strictly entails $\mathcal{B}$. Since $\mathsf{Orep}(\mathcal{A}, \alpha)$ covers all repairs, there is an element $\mathcal{C}$ of $\mathsf{Orep}(\mathcal{A}, \alpha)$ that entails $\mathcal{B}'$, and thus strictly entails $\mathcal{B}$. Consequently, there is $\beta \in \mathcal{C}$ that is not entailed by $\mathcal{B}$. Thus, $\beta \in \mathsf{Cn}^*(\mathcal{A}, \alpha)$, but $\beta \notin \mathcal{B}$, which shows that $\mathcal{B} \subset \mathcal{B} \cup \{\beta\} \subseteq \mathsf{Cn}^*(\mathcal{A}, \alpha)$. This yields a contradiction to our assumption that $\mathcal{B}$ is a remainder of $\mathsf{Cn}^*(\mathcal{A}, \alpha)$ with respect to α if we can show that $\alpha \notin \mathsf{Cn}(\mathcal{B} \cup \{\beta\})$. This finishes the proof since $\alpha \notin \mathsf{Cn}(\mathcal{B} \cup \{\beta\})$ is an easy consequence of the facts that $\mathcal{B} \subseteq \mathsf{Cn}(\mathcal{C})$, $\beta \in \mathcal{C}$, and $\alpha \notin \mathsf{Cn}(\mathcal{C})$. □

As a consequence of this lemma, we can show that maxichoice Cn^* partial meet pseudo-contractions (i.e., ones where the selection function returns a singleton set) always produce optimal repairs in case $\mathsf{Orep}(\mathcal{A}, \alpha)$ covers all repairs.

THEOREM 15. *Let $\mathcal{A}$ be a knowledge base and α a sentence such that $\mathsf{Orep}(\mathcal{A}, \alpha)$ covers all repairs of $\mathcal{A}$ with respect to α. If $\mathsf{ctr}^*_{\gamma^*}$ is a maxichoice Cn^* partial meet pseudo-contraction, then $\mathsf{ctr}^*_{\gamma^*}(\mathcal{A}, \alpha)$ is an optimal repair of $\mathcal{A}$ with respect to α.*

PROOF. In the maxichoice case, the selection function returns a remainder of $\mathsf{Cn}^*(\mathcal{A}, \alpha)$ with respect to α. By Lemma 14, this remainder is an optimal repair of $\mathcal{A}$ with respect to α. □

Since every kernel contraction induced by a minimal incision function can be obtained as a maxichoice partial meet contraction, the theorem also holds if we replace "maxichoice Cn^* partial meet pseudo-contraction" with "Cn^* kernel pseudo-contraction induced by a minimal incision function."

COROLLARY 16. *Let $\mathcal{A}$ be a knowledge base and α a sentence such that $\mathsf{Orep}(\mathcal{A}, \alpha)$ covers all repairs of $\mathcal{A}$ with respect to α. If $\mathsf{ctr}^*_{\delta^*}$ is a Cn^* kernel pseudo-contraction induced by a minimal incision function δ^*, then $\mathsf{ctr}^*_{\delta^*}(\mathcal{A}, \alpha)$ is an optimal repair of $\mathcal{A}$ with respect to α.*

In the ABox repair setting, the condition that $\mathsf{Orep}(\mathcal{A}, \alpha)$ covers all repairs of $\mathcal{A}$ with respect to α is satisfied if we restrict the ABox to being acyclic and the TBox to being cycle-restricted. The ABox $\mathcal{A}$ is called *cyclic* if, for some $n \geq 1$, there are role names $r_1, \ldots, r_n$ and individual names $a_0, a_1, \ldots, a_n$ such that the role assertions $r_1(a_0, a_1), \ldots, r_n(a_{n-1}, a_n)$ belong to $\mathcal{A}$ and $a_0 = a_n$. Otherwise, $\mathcal{A}$ is called *acyclic*. The $\mathcal{EL}$ TBox $\mathcal{T}$ is called *cycle-restricted* if there is no $\mathcal{EL}$ concept C such that $C \sqsubseteq^{\mathcal{T}} \exists r_1. \cdots \exists r_k.C$ for $k \geq 1$ and role names $r_1, \ldots, r_k$.

PROPOSITION 17 ([6], COROLLARY 20). *If $\mathcal{A}$ is acyclic and $\mathcal{T}$ is cycle-restricted, then $\mathsf{Orep}(\mathcal{A}, \alpha)$ covers all repairs of $\mathcal{A}$ w.r.t. α.*

We have already seen in Example 9 that the proposition need not hold if the TBox is not cycle-restricted. The following example demonstrates why acyclicity of $\mathcal{A}$ is needed.

EXAMPLE 18. Assume that $\mathcal{T} = \emptyset$ and consider the cyclic ABox $\mathcal{A} := \{A(a), r(a, a)\}$. If we set $\alpha := \exists r. A(a)$, then $\mathcal{B} := \{A(a)\}$ is a repair of $\mathcal{A}$ with respect to α. Assume that $\mathcal{C}$ is an optimal repair of $\mathcal{A}$ with respect to α that entails $\mathcal{B}$. Then $\mathcal{C}$ cannot contain the role assertion $r(a, a)$, and thus it can entail $(\exists r.)^n\top(a)$ only for finitely many n. Hence, there is an n such that $\mathcal{C}$ does not entail $(\exists r.)^n\top(a)$, which implies that $\mathcal{C} \cup \{(\exists r.)^n\top(a)\}$ is a repair that strictly entails $\mathcal{C}$. This contradicts our assumption that $\mathcal{C}$ is optimal.

5. CONCEPT CONTRACTIONS AND OPTIMAL REPAIRS

In [29], a concept contraction operation for the description logic $\mathcal{EL}$ has been introduced, which is akin to partial meet contraction, but uses the least common subsumer (lcs) to combine the remainders chosen by the selection function. It is then proved that the contraction operations obtained this way can be characterized by appropriate postulates. In this section, we recall this contraction approach and then show how results obtained in the context of optimal repairs can be used to improve on the results in [29]. Following [29], we start with a setting where the TBox is empty.

5.1 The Case of an Empty TBox

A *concept contraction operation* ctr for $\mathcal{EL}$ accepts $\mathcal{EL}$ concepts C, D as input, and produces as output an $\mathcal{EL}$ concept $\mathsf{ctr}(C, D)$ that satisfies the following two postulates:

- $C \sqsubseteq^\emptyset \mathsf{ctr}(C, D)$ (inclusion),
- if $\top \not\equiv^\emptyset D$, then $\mathsf{ctr}(C, D) \not\sqsubseteq^\emptyset D$ (success).

The interesting case is of course the one where $C \sqsubseteq^\emptyset D$. In fact, if $C \not\sqsubseteq^\emptyset D$, then C itself can be used as contraction. The following postulates restricts the attention to contraction operations that actually choose C in this case:

- if $C \not\sqsubseteq^\emptyset D$, then $\mathsf{ctr}(C, D) \equiv^\emptyset C$ (vacuity).

Another reasonable requirement is that the result of contracting w.r.t. D should depend only on the semantics of D, and not on its syntactic form. This is expressed by the following postulate:

- if $D \equiv^\emptyset D'$, then $\mathsf{ctr}(C, D) \equiv^\emptyset \mathsf{ctr}(C, D')$ (preservation).

Finally, a concept contraction operation should always yield a result, even if there is no concept C' such that $C \sqsubseteq^\emptyset C'$ and $C' \not\sqsubseteq^\emptyset D$. This is obviously the case iff $D \equiv^\emptyset \top$. The following postulate states that, in this case, the input concept C should be returned:

- if $D \equiv^\emptyset \top$, then $\mathsf{ctr}(C, D) \equiv^\emptyset C$ (failure).

This corresponds to the fact that, for the case of belief base contraction, the contraction operations $\mathsf{ctr}(\mathcal{B}, \alpha)$ returns $\mathcal{B}$ if $\alpha \in \mathsf{Cn}(\emptyset)$. However, in the belief base case, no extra postulate is needed since this follows both from *relevance* and from *core-retainment*. For partial meet contractions, this property is achieved by requiring that the selection function returns $\{\mathcal{B}\}$ if the set of remainders is empty.

To construct concept contractions that satisfy the above postulates, the authors of [29] adapt the partial meet contraction approach to the concept case. First, the definitions of remainders and selection functions are transferred from belief bases $\mathcal{B}$ and sentences α to $\mathcal{EL}$ concepts C and D as follows:

- a *remainder* of C with respect to D is an $\mathcal{EL}$ concept C' such that $C \sqsubseteq^\emptyset C'$, $C' \not\sqsubseteq^\emptyset D$, and C' is most specific with this property, i.e., there is no $\mathcal{EL}$ concept C'' such that $C \sqsubseteq^\emptyset C''$, $C'' \not\sqsubseteq^\emptyset D$, and $C'' \sqsubset^\emptyset C'$. As before, we denote the set of all remainders as $\mathsf{rem}(C, D)$.[3]
- A *selection function* γ for C takes such sets of remainders as input and satisfies the following properties for each $\mathcal{EL}$ concept D:
  - $\emptyset \neq \gamma(\mathsf{rem}(C, D)) \subseteq \mathsf{rem}(C, D)$ if $\mathsf{rem}(C, D) \neq \emptyset$,
  - $\gamma(\mathsf{rem}(C, D)) = \{C\}$ if $\mathsf{rem}(C, D) = \emptyset$.

In the case of belief base contraction, the remainders chosen by the selection function are intersected to obtain the partial meet contraction, i.e., this contraction requires only properties that all elements of $\gamma(\mathsf{rem}(\mathcal{B}, \alpha))$ have in common. For concepts rather than belief bases, this corresponds to building the least common subsumer:

- the $\mathcal{EL}$ concept C is a *least common subsumer (lcs)* of the $\mathcal{EL}$ concepts $C_1, \ldots, C_n$ if $C_i \sqsubseteq^\emptyset C$ for $i = 1, \ldots, n$, and C is most specific with this property, i.e. there is no $\mathcal{EL}$ concept C' such that $C_i \sqsubseteq^\emptyset C'$ for $i = 1, \ldots, n$ and $C' \sqsubset^\emptyset C$.

By definition, least common subsumers are unique up to equivalence. In the following, we write $\mathsf{lcs}(C_1, \ldots, C_n)$ to denote an arbitrary element of the equivalence class of the least common subsumers of $C_1, \ldots, C_n$. It was shown in [9] that the lcs of $\mathcal{EL}$ concepts always exists, and can effectively by computed, but its size may be exponential in the size of the input $C_1, \ldots, C_n$, unless n is assumed to be constant.

[3]More precisely, $\mathsf{rem}(C, D)$ contains one representative for each equivalence class of concepts that are remainders.

Given a finite set of $\mathcal{EL}$ concepts, its lcs is the lcs of the sequence of its elements, enumerated in an arbitrary order.

Each selection function γ induces an *lcs partial meet concept contraction* operation ctr_γ as follows:

$$\mathsf{ctr}_\gamma(C, D) := \mathsf{lcs}(\gamma(\mathsf{rem}(C, D))).$$

In [29], both arbitrary lcs partial meet concept contraction operations and *maxichoice* lcs partial meet concept contraction operations, where the selection function always returns a singleton set, are characterized by appropriate postulates. The characterization of the maxichoice variant requires the following additional postulate:

- if $C \sqsubseteq^\emptyset X$ and $\mathsf{ctr}(C, D) \not\sqsubseteq^\emptyset X$, then $\mathsf{ctr}(C, D) \sqcap X \sqsubseteq^\emptyset D$ (fullness).

It is shown in [29] (Theorem 1 in [29]) that a concept contraction operator ctr is a maxichoice lcs partial meet concept contraction iff it satisfies the postulates *preservation*, *inclusion*, *success*, *failure*, and *fullness*. Note that *vacuity* need not be required explicitly since it is implied by *inclusion* and *fullness* (see Proposition 2 in [29]).

In the characterization of arbitrary lcs partial meet concept contraction operations (Theorem 2 in [29]), *fullness* is replaced with *relevance*:

- if $C \sqsubseteq^\emptyset X$ and $\mathsf{ctr}(C, D) \not\sqsubseteq^\emptyset X$, then there is Y s.t. $C \sqsubset^\emptyset Y \sqsubseteq^\emptyset \mathsf{ctr}(C, D)$, $Y \not\sqsubseteq^\emptyset D$, and $Y \sqcap X \sqsubseteq^\emptyset D$ (relevance).

Obviously, *fullness* implies *relevance*.[4] Again, *vacuity* need not be required explicitly since it is implied by *inclusion* and *relevance*.

PROPOSITION 19. *If* ctr *satisfies* inclusion *and* relevance*, then it satisfies* vacuity*.*

PROOF. Assume that ctr satisfies *inclusion* and *relevance*. To show *vacuity*, we consider the situation where $C \not\sqsubseteq^\emptyset D$. We must show that $\mathsf{ctr}(C, D) \equiv^\emptyset C$. Since *inclusion* yields $C \sqsubseteq^\emptyset \mathsf{ctr}(C, D)$, it is enough to prove $\mathsf{ctr}(C, D) \sqsubseteq^\emptyset C$. Assume to the contrary that $\mathsf{ctr}(C, D) \not\sqsubseteq^\emptyset C$. If we set $X := C$, then the prerequisite for *relevance* is satisfied. This yields a concept Y such that $C \sqsubset^\emptyset Y \sqsubseteq^\emptyset \mathsf{ctr}(C, D)$ and $Y \sqcap C \sqsubseteq^\emptyset D$. However, $C \sqsubset^\emptyset Y$ implies $C \equiv^\emptyset Y \sqcap C$, and thus $C \equiv^\emptyset Y \sqcap C \sqsubseteq^\emptyset D$, which contradicts our assumption that $C \not\sqsubseteq^\emptyset D$. □

The proof of Theorem 2 in [29] makes use of the following covering property of remainders.

PROPOSITION 20 ([29]). *If $C \sqsubseteq^\emptyset X$ and $X \not\sqsubseteq^\emptyset D$, then there is a remainder Z of C w.r.t. D such that $Z \sqsubseteq^\emptyset X$.*

If one wants to apply the lcs partial meet concept contraction approach in practice, one needs to be able to compute remainders. Although the main focus of the paper [29] is on proving characterization theorems, the authors also sketch how remainders can in principle be found. For this purpose, they use what they call *most specific generalizations (MSGs)* of $\mathcal{EL}$ concepts. Under the name of upper neighbors, MSGs had been defined and characterized before in [8]: given two $\mathcal{EL}$ concepts C, D, the concept D is an *upper neighbor* of C if $C \sqsubset^\emptyset D$ and there is no $\mathcal{EL}$ concept D' such that $C \sqsubset^\emptyset D' \sqsubset^\emptyset D$. It is shown in [8] that

- the set of upper neighbors of a given $\mathcal{EL}$ concept C contains polynomially many elements (in the size of C), and it can be computed in polynomial time;
- the subsumption pre-order on $\mathcal{EL}$ concepts is *one-step generated*, i.e., if $C \sqsubseteq^\emptyset D$ for $\mathcal{EL}$ concepts C and D, then there are an integer $n \geq 0$ and $\mathcal{EL}$ concepts $C = C_0, \ldots, C_n = D$ such that C_{i+1} is an upper neighbor of C_i for all $i, 0 \leq i < n$;
- the upper neighbor relation is well-founded, i.e., there is no infinite sequence of $\mathcal{EL}$ concepts $C_0, C_1, C_2, \ldots$ such that C_{i+1} is an upper neighbor of C_i for all $i \geq 0$.

Let C, D be $\mathcal{EL}$ concepts such that $C \sqsubseteq^\emptyset D$ and $D \not\equiv^\emptyset \top$.[5] Based on the results for upper neighbors recalled above, one can search for remainders of C w.r.t. D by starting with $C_0 = C$ and generating an upper neighbor sequence $C_0, C_1, C_2, \ldots$ until a concept C_n with $C_n \not\sqsubseteq^\emptyset D$ is reached. There are, however, two problems with this approach. First, while well-foundedness implies that such a concept C_n is always found after finitely many steps, the length of upper neighbor sequences starting with C cannot be bounded by an elementary function in the size of C. In fact, as shown in [22], the length of any upper neighbor sequence from $\exists r_1.\exists r_2.\cdots\exists r_n.(A_1 \sqcap \ldots \sqcap A_k)$ (with $k \geq 3$) to $\top$ is asymptotically bounded from below by a tower of exponential of height n. It is, however, not clear whether such a worst-case example also exists for sequences where the concepts C_i are subsumed by a given $\mathcal{EL}$ concept $D \not\equiv^\emptyset \top$. Second, while the concept C_n found this way satisfies $C \sqsubseteq^\emptyset C_n$ and $C_n \not\sqsubseteq^\emptyset D$, it need not be most specific with this property, as illustrated by the following example.

EXAMPLE 21. Let $C := A_1 \sqcap A_2$ and $D := A_1$. If we consider the upper neighbor sequence $C_0 = C, C_1 = A_1, C_2 = \top$, then C_2 is the first element that is not subsumed by D. For the sequence $C'_0 = C, C'_1 = A_2, C'_2 = \top$, C'_1 is the first element not subsumed by D. The concept $C'_1 = A_2$ is a remainder of C w.r.t. D, whereas $C_2 = \top$ is not.

This second problem demonstrates that, even if one is only interested in finding a single remainder, it is not sufficient to generate just a single upper neighbor sequence starting with C since one cannot be sure that the first concept C_n in this sequence satisfying $C_n \not\sqsubseteq^\emptyset D$ is really a remainder.

Fortunately, instead of using such a blind search for remainders along upper neighbor sequences, as suggested in [29], one can employ an algorithm for computing optimal repairs

[4]In [29], the implication is stated erroneously in the other direction.

[5]If one of these two properties is not satisfied, then C is the only remainder or there is no remainder.

to obtain the set of all remainders. The following lemma, which states the connection between optimal repairs and remainders, is easy to show.

LEMMA 22. *Let C, D be $\mathcal{EL}$ concepts. Then C' is a remainder of C w.r.t. D iff the ABox $\{C'(a)\}$ is an optimal repair of the ABox $\{C(a)\}$ with respect to $\alpha = D(a)$ for the consequence relation $\mathsf{Cn}_\emptyset$.*

The problem of computing optimal repairs for ABoxes consisting of single concept assertions (called instance stores [21]) has been investigated in [7][6] in the more general setting where more than one unwanted consequence can be specified. In terms of concept contraction, this would correspond to the setting where one is given $\mathcal{EL}$ concepts $C, D_1, \ldots, D_n$ and looks for concepts C' such $C \sqsubseteq^\emptyset C'$ and $C' \not\sqsubseteq^\emptyset D_i$ for $i = 1, \ldots, n$. It is shown in [7] that, in the general setting, the set of all optimal repairs can be computed in exponential time. Up to equivalence, there are at most exponentially many optimal repairs of at most exponential size. Example 2 in [7] demonstrates that these exponential bounds are tight. However, for the case of a single unwanted consequence, a close inspection of Definition 4 in [7] reveals that in this special case there are only linearly many optimal repairs of polynomial size, which can be computed in polynomial time. Together with Lemma 22, this yields the following complexity result for computing remainders.

THEOREM 23. *Let C, D be $\mathcal{EL}$ concepts. The set of all remainders of C w.r.t. D can be computed in polynomial time.*

Since lcs partial meet concept contraction operations return the lcs of a subset of the set of all remainders, and the size of such an lcs may be exponential in the cardinality of this set [9], we obtain the following complexity results for computing such concept contractions.

COROLLARY 24. *If the selection function γ can be computed in exponential time, then the lcs partial meet concept contraction operation ctr_γ can be computed in exponential time. In addition, if the selection function γ always returns singleton sets and can be computed in polynomial time, then the maxichoice lcs partial meet concept contraction operation ctr_γ can be computed in polynomial time.*

5.2 The Case of a Non-Empty TBox

Given an $\mathcal{EL}$ TBox $\mathcal{T}$, the definitions and postulates introduced in the previous section can be reformulated into $\mathcal{T}$-variants, by replacing subsumption w.r.t. the empty TBox ($\sqsubseteq^\emptyset$) with subsumption w.r.t. $\mathcal{T}$ ($\sqsubseteq^\mathcal{T}$).

For instance, a *$\mathcal{T}$-remainder* of C w.r.t. D is an $\mathcal{EL}$ concept C' such that $C \sqsubseteq^\mathcal{T} C'$, $C' \not\sqsubseteq^\mathcal{T} D$, and C' is most specific with this property, i.e., there is no $\mathcal{EL}$ concept C'' such that $C \sqsubseteq^\mathcal{T} C''$, $C'' \not\sqsubseteq^\mathcal{T} D$, and $C'' \sqsubset^\mathcal{T} C'$; the $\mathcal{T}$-variant of the postulate *success* is

- if $\top \not\equiv^\mathcal{T} D$, then $\mathsf{ctr}(C, D) \not\sqsubseteq^\mathcal{T} D$ ($\mathcal{T}$-success);

[6]More precisely, optimal repairs are called optimal compliant generalizations in [7].

and the lcs w.r.t. $\mathcal{T}$ is defined as follows:

- the $\mathcal{EL}$ concept C is a *least common subsumer ($\mathcal{T}$-lcs)* of the $\mathcal{EL}$ concepts $C_1, \ldots, C_n$ *w.r.t.$\mathcal{T}$* if $C_i \sqsubseteq^\mathcal{T} C$ for $i = 1, \ldots, n$, and C is most specific with this property, i.e. there is no $\mathcal{EL}$ concept C' such that $C_i \sqsubseteq^\mathcal{T} C'$ for $i = 1, \ldots, n$ and $C' \sqsubset^\mathcal{T} C$.

The definitions of concept contraction operations and (maxichoice) lcs partial meet concept contraction operations can then be adapted to the setting of a non-empty TBox in the obvious way. However, for these definitions to produce sensible results, the existence of $\mathcal{T}$-remainders and of the $\mathcal{T}$-lcs needs to be guaranteed

Without restrictions on the TBox $\mathcal{T}$, this is neither the case for $\mathcal{T}$-remainders nor for the $\mathcal{T}$-lcs. For $\mathcal{T}$-remainders, this is shown in [29] with an example that is basically the same as our Example 8 demonstrating non-existence of optimal repairs. For the $\mathcal{T}$-lcs, this is, e.g., shown in [2]. To overcome this problem, the authors of [29] restrict the attention to acyclic TBoxes, which basically only introduce concept names (called defined concepts) as abbreviations for compound concepts [4]. For a given acyclic TBox $\mathcal{T}$, concepts C containing defined concepts can be expanded into concepts $C^\mathcal{T}$ not containing defined concepts by iteratedly replacing defined concepts by their definitions. The idea underlying the concept contraction approach for acyclic TBoxes introduced in [29] is now to first expand the input concepts C and D w.r.t. $\mathcal{T}$, and then apply the concept contraction approach w.r.t. the empty TBox to the expanded concepts.

In the following, we show how to deal with a considerably larger class of TBoxes, which are the *cycle-restricted* TBoxes introduced above Proposition 17 in Section 4. First note that Lemma 22 obviously also holds for the case of a non-empty TBox.

LEMMA 25. *Let C, D be $\mathcal{EL}$ concepts and $\mathcal{T}$ an $\mathcal{EL}$ TBox. Then C' is a $\mathcal{T}$-remainder of C w.r.t. D iff the ABox $\{C'(a)\}$ is an optimal repair of the ABox $\{C(a)\}$ with respect to $\alpha = D(a)$ for the consequence relation $\mathsf{Cn}_\mathcal{T}$.*

If $\mathcal{T}$ is cycle-restricted, then Proposition 17 implies that these optimal repairs actually cover all repairs. This shows that Proposition 20 also holds w.r.t. such TBoxes rather than just w.r.t. the empty TBox.

PROPOSITION 26. *Let $\mathcal{T}$ be a cycle-restricted TBox. If $C \sqsubseteq^\mathcal{T} X$ and $X \not\sqsubseteq^\mathcal{T} D$, then there is a $\mathcal{T}$-remainder Z of C w.r.t. D such that $Z \sqsubseteq^\mathcal{T} X$*

Second, the characterization of the existence of the $\mathcal{T}$-lcs given in [33] implies that the $\mathcal{T}$-lcs always exists for cycle-restricted TBoxes $\mathcal{T}$. Thus, the definition of lcs partial meet concept contraction operations also makes w.r.t. such TBoxes. With Proposition 26 in place, it is now easy to check that the proofs of Theorem 1 and 2 in [29] also go through for cycle-restricted TBoxes.

THEOREM 27. *Let $\mathcal{T}$ be a cycle-restricted TBox. Then*

1. ctr *is a maxichoice lcs concept contraction operation w.r.t. $\mathcal{T}$ iff it satisfies* $\mathcal{T}$-preservation*,* $\mathcal{T}$-inclusion*,* $\mathcal{T}$-success*,* $\mathcal{T}$-failure*, and* $\mathcal{T}$-fullness.

2. ctr *is an lcs concept contraction operation w.r.t.* $\mathcal{T}$ *iff it satisfies* $\mathcal{T}$-preservation*,* $\mathcal{T}$-inclusion*,* $\mathcal{T}$-success*,* $\mathcal{T}$-failure*, and* $\mathcal{T}$-relevance*.*

Finally, for cycle-restricted TBoxes, $\mathcal{T}$-remainders and the $\mathcal{T}$-lcs not only exist, but can also effectively be computed. The complexity result for computing optimal ABox repairs proved in [6] yields the following complexity upper bound for computing $\mathcal{T}$-remainders.

THEOREM 28. *Let* C, D *be* $\mathcal{EL}$ *concepts and* $\mathcal{T}$ *a cycle-restricted* $\mathcal{EL}$ *TBox. The set of all* $\mathcal{T}$*-remainders of* C *w.r.t.* D *can be computed in double-exponential time.*

It should be noted that the double-exponential upper bound is shown in [6] for a considerably more general setting, and thus the complexity in our restricted setting might actually be lower. Since computability of the $\mathcal{T}$-lcs is shown in [33], we can also compute the (maxichoice) lcs concept contraction operations w.r.t. $\mathcal{T}$.

COROLLARY 29. *Let* $\mathcal{T}$ *be a cycle-restricted* $\mathcal{EL}$ *TBox. If the selection function* γ *is computable, then the lcs partial meet concept contraction operation* ctr_γ *w.r.t.* $\mathcal{T}$ *is also computable.*

6. CONCLUSION

The results shown is this paper complement recent results [24, 25] on the relationship between gentle repairs and pseudo-contractions by demonstrating that there are close connections between optimal repairs and certain pseudo-contraction operations. We have illustrated these results on the use case of repairing $\mathcal{EL}$ ABoxes with respect to static $\mathcal{EL}$ TBoxes, where optimal repairs can effectively be computed (if they exists) [6].

In [5], it was shown that optimal repairs always exist and cover all repairs if one uses quantified ABoxes (where some of the individuals can be anonymized by representing them as existentially quantified variables) in place of ABoxes. Extending the result of the present paper to this setting poses new challenges since the first-order translation of a quantified ABox is not a set of sentences, but a single one, which starts with an existential quantifier prefix. Thus, considering subsets when constructing contractions does not make sense. We conjecture that this problem can be overcome by introducing an "inclusion" relation on quantified ABoxes that shares enough properties with set inclusion for the constructions and proofs regarding (pseudo-)contractions to continue working.

On a more conceptual level, there are certain differences between repair approaches in ontology engineering and contraction approaches in belief change that are worth investigating. On the one hand, the work on optimal repairs [5, 6] usually considers a single repair problem and does not investigate the relationship between repairs for different unwanted consequences, whereas postulates like *uniformity* in belief change make statements on how results for different unwanted consequences should be connected under certain conditions on these consequences. It would be interesting to see whether and how postulates like uniformity and their variants in the context of pseudo-contractions [30, 25] can be satisfied by methods that compute optimal repairs. On the other hand, contraction and pseudo-contraction operators produces a single belief base as output, whereas work on optimal repairs is also concerned with how to compute the set of all such repairs and investigates properties of this set (like whether it covers all repairs or not). In contrast, on the belief change side, there are no postulates about the sets of all pseudo-contractions that can be obtained be applying a certain approach (e.g., in the partial meet case, if one looks at all possible selection functions). It would be interesting to see whether taking this "set view" can lead to interesting kinds of new postulates.

Regarding concept contractions, it remains to determine the exact complexity of computing remainders and lcs partial meet concept contraction operations for the case of cycle-restricted TBoxes. On a more conceptual level, it would be interesting to see whether one can also adapt the kernel contraction approach to this setting.

7. ACKNOWLEDGEMENTS

This work was partially supported by the German Research Foundation (DFG) in Project 430150274 and SFB/TRR 248. The author should like to thank the anonymous reviewers of the conference SAC 2023 for their helpful comments, and Francesco Kriegel and Adrian Nuradiansyah for helpful discussions.

8. REFERENCES

[1] C. E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *J. Symb. Log.*, 50(2):510–530, 1985.

[2] F. Baader. Computing the least common subsumer in the description logic EL w.r.t. terminological cycles with descriptive semantics. In A. de Moor, W. Lex, and B. Ganter, editors, *Conceptual Structures for Knowledge Creation and Communication, 11th International Conference on Conceptual Structures, ICCS 2003 Dresden, Germany, July 21-25, 2003 Proceedings*, volume 2746 of *Lecture Notes in Computer Science*, pages 117–130. Springer, 2003.

[3] F. Baader. Optimal repairs in ontology engineering as pseudo-contractions in belief change. In J. Hong, M. Lanperne, J. W. Park, T. Cerný, and H. Shahriar, editors, *Proceedings of the 38th ACM/SIGAPP Symposium on Applied Computing, SAC 2023, Tallinn, Estonia, March 27-31, 2023*, pages 983–990. ACM, 2023.

[4] F. Baader, I. Horrocks, C. Lutz, and U. Sattler. *An Introduction to Description Logic.* Cambridge University Press, 2017.

[5] F. Baader, P. Koopmann, F. Kriegel, and A. Nuradiansyah. Computing optimal repairs of quantified ABoxes w.r.t. static $\mathcal{EL}$ TBoxes. In A. Platzer and G. Sutcliffe, editors, *Automated Deduction - CADE 28 - 28th International Conference on Automated Deduction, Proceedings*, volume 12699 of *LNCS*, pages 309–326. Springer, 2021.

[6] F. Baader, P. Koopmann, F. Kriegel, and A. Nuradiansyah. Optimal ABox repair w.r.t. static $\mathcal{EL}$ TBoxes: From quantified ABoxes back to ABoxes. In *The Semantic Web - 19th International Conference, ESWC 2022, Proceedings*, volume 13261 of *LNCS*, pages 130–146. Springer, 2022.

[7] F. Baader, F. Kriegel, and A. Nuradiansyah. Privacy-preserving ontology publishing for $\mathcal{EL}$ instance stores. In *Logics in Artificial Intelligence - 16th European Conference, JELIA 2019, Proceedings*, volume 11468 of *LNCS*. Springer, 2019.

[8] F. Baader, F. Kriegel, A. Nuradiansyah, and R. Peñaloza. Making repairs in description logics more gentle. In M. Thielscher, F. Toni, and F. Wolter, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018, Tempe, Arizona, 30 October - 2 November 2018*, pages 319–328. AAAI Press, 2018.

[9] F. Baader, R. Küsters, and R. Molitor. Computing least common subsumers in description logics with existential restrictions. In T. Dean, editor, *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence, IJCAI 99, Stockholm, Sweden, July 31 - August 6, 1999. 2 Volumes, 1450 pages*, pages 96–103. Morgan Kaufmann, 1999.

[10] F. Baader and B. Suntisrivaraporn. Debugging SNOMED CT using axiom pinpointing in the description logic $\mathcal{EL}^+$. In R. Cornet and K. A. Spackman, editors, *Proceedings of the Third International Conference on Knowledge Representation in Medicine, Phoenix, Arizona, USA, May 31st - June 2nd, 2008*, volume 410 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2008.

[11] J. Du, G. Qi, and X. Fu. A practical fine-grained approach to resolving incoherent OWL 2 DL terminologies. In *Proc. of the 23rd ACM Int. Conf. on Information and Knowledge Management, (CIKM'14)*, pages 919–928. ACM, 2014.

[12] M. A. Falappa, E. L. Fermé, and G. Kern-Isberner. On the logic of theory change: Relations between incision and selection functions. In G. Brewka, S. Coradeschi, A. Perini, and P. Traverso, editors, *ECAI 2006, 17th European Conference on Artificial Intelligence, August 29 - September 1, 2006, Riva del Garda, Italy, Including Prestigious Applications of Intelligent Systems (PAIS 2006), Proceedings*, volume 141 of *Frontiers in Artificial Intelligence and Applications*, pages 402–406. IOS Press, 2006.

[13] A. Fuhrmann. Theory contraction through base contraction. *J. Philos. Logic*, 20:175–203, 1991.

[14] P. Gärdenfors, editor. *Belief Revision*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 1992.

[15] S. O. Hansson. New operators for theory change. *Theoria*, 55(2):114–132, 1989.

[16] S. O. Hansson. In defense of base contraction. *Synthese*, 92:239–245, 1992.

[17] S. O. Hansson. Changes of disjunctively closed bases. *J. Log. Lang. Inf.*, 2(4):255–284, 1993.

[18] S. O. Hansson. Reversing the Levi identity. *J. Philos. Logic*, 22(6):637–669, 1993.

[19] S. O. Hansson. Kernel contraction. *J. Symb. Log.*, 59(3):845–859, 1994.

[20] M. Horridge, B. Parsia, and U. Sattler. Laconic and precise justifications in OWL. In A. P. Sheth, S. Staab, M. Dean, M. Paolucci, D. Maynard, T. W. Finin, and K. Thirunarayan, editors, *The Semantic Web - ISWC 2008, 7th International Semantic Web Conference, ISWC 2008, Karlsruhe, Germany, October 26-30, 2008. Proceedings*, volume 5318 of *Lecture Notes in Computer Science*, pages 323–338. Springer, 2008.

[21] I. Horrocks, L. Li, D. Turi, and S. Bechhofer. The instance store: DL reasoning with large numbers of individuals. In V. Haarslev and R. Möller, editors, *Proceedings of the 2004 International Workshop on Description Logics (DL2004), Whistler, British Columbia, Canada, June 6-8, 2004*, volume 104 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2004.

[22] F. Kriegel. Navigating the $\mathcal{EL}$ subsumption hierarchy. In M. Homola, V. Ryzhikov, and R. A. Schmidt, editors, *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*, volume 2954 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2021.

[23] J. S. C. Lam, D. H. Sleeman, J. Z. Pan, and W. W. Vasconcelos. A fine-grained approach to resolving unsatisfiable ontologies. *J. Data Semant.*, 10:62–95, 2008.

[24] V. B. Matos, R. Guimarães, Y. D. Santos, and R. Wassermann. Pseudo-contractions as gentle repairs. In C. Lutz, U. Sattler, C. Tinelli, A. Turhan, and F. Wolter, editors, *Description Logic, Theory Combination, and All That - Essays Dedicated to Franz Baader on the Occasion of His 60th Birthday*, volume 11560 of *Lecture Notes in Computer Science*, pages 385–403. Springer, 2019.

[25] V. B. Matos and R. Wassermann. Repairing ontologies via kernel pseudo-contraction. In O. Arieli, G. Casini, and L. Giordano, editors, *Proceedings of the 20th International Workshop on Non-Monotonic Reasoning, NMR 2022, Part of the Federated Logic Conference (FLoC 2022), Haifa, Israel, August 7-9, 2022*, volume 3197 of *CEUR Workshop Proceedings*, pages 16–26. CEUR-WS.org, 2022.

[26] B. Parsia, E. Sirin, and A. Kalyanpur. Debugging OWL ontologies. In A. Ellis and T. Hagino, editors, *Proceedings of the 14th International Conference on World Wide Web, WWW 2005, Chiba, Japan, May 10-14, 2005*, pages 633–640. ACM, 2005.

[27] G. Qi and F. Yang. A survey of revision approaches in description logics. In D. Calvanese and G. Lausen, editors, *Web Reasoning and Rule Systems, Second International Conference, RR 2008, Karlsruhe, Germany, October 31-November 1, 2008. Proceedings*, volume 5341 of *Lecture Notes in Computer Science*, pages 74–88. Springer, 2008.

[28] R. Reiter. A theory of diagnosis from first principles. *Artif. Intell.*, 32(1):57–95, 1987.

[29] T. Rienstra, C. Schon, and S. Staab. Concept contraction in the description logic $\mathcal{EL}$. In D. Calvanese, E. Erdem, and M. Thielscher, editors, *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning, KR 2020, Rhodes, Greece, September 12-18, 2020*, pages 723–732, 2020.

[30] Y. D. Santos, V. B. Matos, M. M. Ribeiro, and R. Wassermann. Partial meet pseudo-contractions. *Int. J. Approx. Reason.*, 103:11–27, 2018.

[31] S. Schlobach, Z. Huang, R. Cornet, and F. Harmelen. Debugging incoherent terminologies. *J. Automated Reasoning*, 39(3):317–349, 2007.

[32] N. Troquard, R. Confalonieri, P. Galliani, R. Peñaloza, D. Porello, and O. Kutz. Repairing ontologies via axiom weakening. In S. A. McIlraith and K. Q. Weinberger, editors, *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18)*, pages 1981–1988. AAAI Press, 2018.

[33] B. Zarrieß and A. Turhan. Most specific generalizations w.r.t. general $\mathcal{EL}$-tboxes. In F. Rossi, editor, *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013*, pages 1191–1197. IJCAI/AAAI, 2013.