Errors Need Not Show Up as Inconsistencies

From Error-tolerant Reasoning to Argumentation Frameworks

Franz Baader[1,2], Oscar Barreca[1,2]

[1] Technische Universität Dresden

[2] Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI) Dresden/Leipzig

Abstract

Description logic (DL) knowledge bases (KBs) built by hand or (semi)automatically using machine learning or information retrieval tools may contain errors, often detected when reasoning derives an inconsistency. However, not all errors cause an inconsistency. Such errors may be noticed when reasoning produces a consequence that follows from the KB, but does not hold in the application domain modelled by the KB. In fact, some application-relevant DL KBs such as the large medical ontology SNOMED CT are written in DL dialects, like $\mathcal{EL}$, that cannot even express inconsistencies. Eliminating an inconsistency or an unwanted consequence by removing a minimal amount of information from the KB is usually called *repair* in the DL literature. Instead of producing a single new KB as a repair, *inconsistency-tolerant reasoning* takes all repairs of a detected inconsistency into account, e.g., by producing consequences that follow from all repairs or from some repair. This notion can be extended to repairs that remove an unwanted consequence, in which case it is called *error-tolerant reasoning*. This paper shows in an exemplary way how results for inconsistency-tolerant reasoning can be transferred to error-tolerant reasoning. To this purpose, it generalizes an existing approach connecting inconsistency-tolerant reasoning over prioritized DL-Lite KBs and argumentation frameworks to the setting of error-tolerant reasoning for $\mathcal{EL}$ KBs.

Keywords

Description logics, abstract argumentation, inconsistency-tolerant reasoning, ontology-based data access

1. Introduction

Inconsistency-tolerant reasoning had originally been considered in logic [1] and for databases [2], but was then extended to the setting of ontology-mediated query answering (see, e.g., [3, 4, 5]). In this context, one usually considers a terminology (TBox) written in some DL and an ABox that contains atomic concept assertions and role assertions. The TBox is assumed to be correct, which means that the inconsistency must be caused by the ABox assertions, in the sense that the ABox together with the TBox does not have a model. A *repair* is then a maximal subset of the given ABox that is consistent with the given TBox. As queries one usually considers conjunctive queries or instance queries (i.e., queries expressible by a DL concept). The question is then whether a certain answer to such a query (i.e., one that logically follows from the TBox together with the ABox) holds for all repairs, for some repair, or for the intersection of all repairs. Most of the work in this direction has been concentrated on investigating the computational complexity of answering this question. Here, we are interested in more foundational work (such as [6, 7, 8, 9, 10]) that links certain forms of inconsistency-tolerant reasoning with the notion of extensions in argumentation frameworks (AFs) [11, 12].

In particular, we consider [9], where the TBox is assumed to be written in common DL-Lite dialects and the ABox is additionally equipped with a priority relation that can be used to resolve conflicts between ABox assertions. The authors define repairs of an inconsistent KB as maximal subsets of the ABox consistent with the TBox, and then use the priority relation to characterize Pareto-optimal repairs. As inconsistency-tolerant semantics they employ (among others) ones that consider reasoning w.r.t. all repairs (AR) and with respect to the intersections of all repairs (IAR), where repairs are either arbitrary repairs or Pareto-optimal repairs. In addition to investigating the complexity of reasoning w.r.t. these semantics, they also try to characterize them using AFs. More precisely, they translate prioritized KBs

Logical approaches to handling INconsistent DAta (LINDA) 2026

franz.baader@tu-dresden.de (F. Baader); oscar.barreca@tu-dresden.de (O. Barreca)

0000-0002-4049-221X (F. Baader); 0009-0005-9257-8217 (O. Barreca)

into PSETAFs, which are preference-based variants of set-based AFs (SETAFs) [13]. Two interesting results obtained from this translation are that Pareto-optimal repairs are exactly the stable extensions of the corresponding PSETAF, and that, in the case of an empty priority relation, the intersection of all repairs coincides with the (unique) grounded extension of the corresponding SETAF.

An inconsistent KB is obviously erroneous, but consistent KBs may also contain errors. For example, in several earlier versions of the medical ontology SNOMED CT,[1] which at this time was a TBox written in the DL $\mathcal{EL}$, the concept "amputation of finger" was classified as a subconcept of "amputation of hand", which should clearly not be the case [14, 15]. If a DL reasoner produces a consequence α that is considered to be unwanted by the user, then the KB must be repaired. In analogy to the repair of an inconsistency this means that one is looking for a maximal subset of the KB that does not have the consequence α.

In [16], such a repair is called an *optimal classical repair*, but to be more in line with the notation used in the inconsistency-tolerant setting, we call it α-repair to clarify what is the consequence that should be removed. Like inconsistency-tolerant reasoning, *error-tolerant reasoning* determines what follows from all α-repairs, some α-repair, or the intersection of all α-repairs. In the context of DLs, error-tolerant reasoning was for the first time introduced in [17], but, motivated by the above example of an error in SNOMED CT, for the case where the TBox needs to be repaired. Error-tolerant reasoning for a setting where the TBox is assumed to be correct and the ABox must be repaired was investigated in [18, 19, 20], but for optimal repairs, which maximize the set of consequences rather than the set of axioms, instead of optimal classical repairs.

The present paper is concerned with error-tolerant reasoning for ABoxes, considering α-repairs for an unwanted consequence α rather than for an inconsistency, and using $\mathcal{EL}$ TBoxes rather than the DL-Lite TBoxes of [9]. Our goal is to show that links between error-tolerant reasoning and certain notions of extensions in AFs, analogous to the ones in [9] mentioned above, can also be established in the case of α-repairs.

2. Background

We introduce KBs of the DL $\mathcal{EL}$ and $\mathcal{ELO}^\bot$, an extension of $\mathcal{EL}$ that allows to express nominals and inconsistencies. Reasoning is tractable in both logics [21].

$\mathcal{ELO}^\bot$ concepts are built inductively, starting with concept names A drawn from a set N_C, and using the concept constructors $\top$ (top concept), $\bot$ (bottom concept), $\{a\}$ (nominal), $C \sqcap D$ (conjunction), and $\exists r.C$ (existential restriction), where C, D are $\mathcal{ELO}^\bot$ concepts, r belongs to a set N_R of role names and a belongs to a set N_I of individual names. An $\mathcal{ELO}^\bot$ *general concept inclusion* (GCI) is an expression of the form $C \sqsubseteq D$ for two $\mathcal{ELO}^\bot$ concepts C, D. An $\mathcal{ELO}^\bot$ TBox is a finite set of GCIs, and an ABox is a finite set of concept assertions $A(a)$ and role assertions $r(a, b)$, where $A \in N_C$, $a, b \in N_I$, and $r \in N_R$. An $\mathcal{EL}$ concept is an $\mathcal{ELO}^\bot$ concept that contains neither $\bot$ nor nominals, and an $\mathcal{EL}$ TBox is an $\mathcal{ELO}^\bot$ TBox whose GCIs are built using only $\mathcal{EL}$ concepts. An $\mathcal{ELO}^\bot$ ($\mathcal{EL}$) KB $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ is a pair composed of an $\mathcal{ELO}^\bot$ ($\mathcal{EL}$) TBox $\mathcal{T}$ and an ABox $\mathcal{A}$.

The semantics of $\mathcal{ELO}^\bot$ is defined in a model-theoretic way, using the notion of an interpretation $\mathcal{I}$, which is a pair $\mathcal{I} = (\Delta^\mathcal{I}, \cdot^\mathcal{I})$, where $\Delta^\mathcal{I}$ is a non-empty set and $\cdot^\mathcal{I}$ is an interpretation function that maps each concept name $A \in N_C$ to a set $A^\mathcal{I} \subseteq \Delta^\mathcal{I}$, each role name $r \in N_R$ to a binary relation $r^\mathcal{I} \subseteq \Delta^\mathcal{I} \times \Delta^\mathcal{I}$, and each individual name $a \in N_I$ to an element $a^\mathcal{I} \in \Delta^\mathcal{I}$. The interpretation of an $\mathcal{ELO}^\bot$ concept is defined inductively as follows: $\top^\mathcal{I} := \Delta^\mathcal{I}$, $\bot^\mathcal{I} := \emptyset$, $\{a\}^\mathcal{I} := \{a^\mathcal{I}\}$, $(C \sqcap D)^\mathcal{I} := C^\mathcal{I} \cap D^\mathcal{I}$, and $(\exists r.C)^\mathcal{I} := \{d \in \Delta^\mathcal{I} \mid \exists e \in C^\mathcal{I} \text{ such that } (d, e) \in r^\mathcal{I}\}$. A model $\mathcal{I}$ of the $\mathcal{ELO}^\bot$ TBox $\mathcal{T}$ is an interpretation that satisfies all its GCIs, i.e., $C^\mathcal{I} \subseteq D^\mathcal{I}$ holds for all $C \sqsubseteq D \in \mathcal{T}$, and a model $\mathcal{I}$ of the ABox $\mathcal{A}$ is an interpretation that satisfies all its assertion, i.e., $a^\mathcal{I} \in A^\mathcal{I}$ for all $A(a) \in \mathcal{A}$ and $(a^\mathcal{I}, b^\mathcal{I}) \in r^\mathcal{I}$ for all $r(a, b) \in \mathcal{A}$. A model of a KB $(\mathcal{T}, \mathcal{A})$ is a model of both $\mathcal{T}$ and $\mathcal{A}$.

Regarding reasoning, for the purpose of this paper we are mainly interested in the instance problem. Given an $\mathcal{ELO}^\bot$ KB $\mathcal{K} = (\mathcal{T}, \mathcal{A})$, an $\mathcal{ELO}^\bot$ concept C, and an individual name a, we say that the

[1] https://www.snomed.org/

concept assertion $\alpha = C(a)$ follows from $\mathcal{K}$ (written $\mathcal{K} \models \alpha$) if every model $\mathcal{I}$ of $\mathcal{K}$ satisfies α, i.e., $a^{\mathcal{I}} \in A^{\mathcal{I}}$ holds. More generally, we will also consider answering Boolean conjunctive queries (BCQ), which can be seen as conjunctions of assertions, which in addition to individual names may contain variables that are existentially quantified. Such a query q is entailed by a KB $\mathcal{K}$ (written $\mathcal{K} \models q$) if every model of $\mathcal{K}$ is a model of q (see [9] for a more detailed definition). To express that $\mathcal{K} \models \alpha$ ($\mathcal{K} \models q$), we will say that $\mathcal{A}$ $\mathcal{T}$-entails α (q). We say that the ABox $\mathcal{A}$ $\mathcal{T}$-entails an inconsistency if the KB $(\mathcal{T}, \mathcal{A})$ does not have a model. Note that this can happen in $\mathcal{ELO}^{\bot}$, but not in $\mathcal{EL}$.

3. Error-Tolerant Semantics over Prioritized KBs

The paper [9] describes how a prioritized KB, a KB where preferences between conflicting assertions are provided by the user, can be translated into an appropriate AF and how the stable and grounded semantics of abstract argumentation [11, 12] can be used to capture two types of repairs of inconsistencies. This section adapts the relevant definitions and results in [9] from inconsistency-tolerant reasoning to error-tolerant reasoning, where the error is an unwanted entailed concept assertion. In [9], TBoxes within KBs are built using certain DL-Lite dialects. Here we assume that we have an $\mathcal{EL}$ KB, i.e., whenever we say KB (without additional qualification) in this section we mean an $\mathcal{EL}$ KB. The more general notion of $\mathcal{ELO}^{\bot}$ KB is only needed in a proof. Since $\mathcal{EL}$ cannot express inconsistency, the notion of inconsistency-tolerant reasoning does not make sense for it, but it does make sense for $\mathcal{ELO}^{\bot}$. We also assume in this section that $\alpha = C(a)$ is an $\mathcal{EL}$ concept assertion, where C is an an $\mathcal{EL}$ concept (which need not be a concept name).

Definition 1 (α-repair). *Given a KB $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ and a concept assertion α, an α-repair is a subset-maximal subset of $\mathcal{A}$ that does not $\mathcal{T}$-entail α. The set of α-repairs for $\mathcal{K}$ is denoted by $Rep_{\alpha}(\mathcal{K})$.*

Compared to the definition of repair in [9], the requirement that a repair must be a $\mathcal{T}$-consistent subset of assertions is replaced by the condition that α-repairs must not $\mathcal{T}$-entail α. The notion of *conflict* is also adapted in a corresponding way. Our new notion of α-conflict is sometimes also called *cause* or *support* for α [22, 23].

Definition 2 (α-conflict). *Given a KB $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ and an assertion α, an α-conflict is a subset-minimal subset of $\mathcal{A}$ that $\mathcal{T}$-entails α. The set of α-conflicts for $\mathcal{K}$ is denoted by $Conf_{\alpha}(\mathcal{K})$.*

The DL $\mathcal{ELO}^{\bot}$ can express inconsistency, e.g., using the assertion $\bot(a)$. Thus, for an $\mathcal{ELO}^{\bot}$ KB, the set $\text{Conf}_{\bot(a)}(\mathcal{K})$ corresponds to the set $\text{Conf}(\mathcal{K})$ of conflicts defined in [9], i.e, the set of subset-minimal subsets of $\mathcal{A}$ $\mathcal{T}$-entailing an inconsistency.

3.1. α-prioritized KBs

In [9], *priority relations* are considered, which are relations between facts establishing a preference between conflicting facts. This notion also needs to be adapted to the case of α-conflicts.

Definition 3 (α-priority relation). *Given a KB $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ and an assertion α, an* α-priority relation *for $\mathcal{K}$ and α is an acyclic relation $\succ_{\alpha} \subseteq \mathcal{A} \times \mathcal{A}$ such that if $\beta \succ_{\alpha} \gamma$, then there exists an α-conflict $C \in Conf_{\alpha}(\mathcal{K})$ such that $\{\beta, \gamma\} \subseteq C$.*

An α-prioritized KB is a KB additionally equipped with such a priority relation.

Definition 4 (α-prioritized KB). *An* α-prioritized *KB is a pair $\mathcal{K}_{\succ_{\alpha}} = (\mathcal{K}, \succ_{\alpha})$ such that $\mathcal{K}$ is a KB and $\succ_{\alpha}$ an α-priority relation for $\mathcal{K}$ and α.*

In practice, the priority relation is usually obtained from a domain expert. Such a relation can be used to select from the set of all repairs certain preferred ones that take the expert knowledge into account.

Definition 5 (Pareto-optimal)**.** *Given an α-prioritized KB $\mathcal{K}_{\succ_\alpha} = (\mathcal{K}, \succ_\alpha)$ for $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ and an α-repair $R \in Rep_\alpha(\mathcal{K})$, a set $S \subseteq \mathcal{A}$ of assertions is said to be a* Pareto improvement *of R with respect to $\succ_\alpha$ if it does not $\mathcal{T}$-entail α and a $\gamma \in S \setminus R$ exists such that $\gamma \succ_\alpha \beta$ for all $\beta \in R \setminus S$. The α-repair R is said to be a* Pareto-optimal α-repair of $\mathcal{K}_{\succ_\alpha}$ *if there is no Pareto improvement of R with respect to $\succ_\alpha$. The set of Pareto-optimal α-repairs of $\mathcal{K}_{\succ_\alpha}$ with respect to $\succ_\alpha$ is denoted by $Prep_\alpha(\mathcal{K}_{\succ_\alpha})$.*

Since the *all repairs* (AR), *intersection of all repairs* (IAR) and *brave* semantics abstract away from the concrete notion of repair, they are defined in [9] in a parameterized away. The same is possible here, except that α-priority relations are considered instead of priority relations.

Definition 6 (Error-tolerant reasoning)**.** *Given an assertion α and a repair type $T \in \{Rep_\alpha, Prep_\alpha\}$, the α-prioritized KB $\mathcal{K}_{\succ_\alpha}$ entails a BCQ or concept assertion q under the:*

- AR semantics for T *(T-AR) if $(\mathcal{T}, R) \models q$ for every $R \in T(\mathcal{K}_{\succ_\alpha})$;*
- IAR semantics for T *(T-IAR) if $(\mathcal{T}, I) \models q$, where $I := \bigcap_{R \in T(\mathcal{K}_{\succ_\alpha})} R$;*
- Brave semantics for T *(T-brave) if $(T, R) \models q$ for some $R \in T(\mathcal{K}_{\succ_\alpha})$.*

It is easy to see that entailment w.r.t. IAR semantics implies entailment w.r.t. AR semantics, which in turn implies entailment w.r.t. brave semantics (see also the remark after Definition 7 in [9]).

3.2. From α-prioritized KBs to AFs

In [9], the authors describe a translation from prioritized KBs to PSETAFs such that the Pareto-optimal repairs and the intersection of all repairs correspond to certain notions of extensions. This translation and the corresponding results can be adapted to our setting of α-repairs.

Due to space constraints, we cannot define PSETAFs and the relevant notions of extensions in detail, and refer thus to [9] for the exact definitions. For what follows, it is sufficient to note that a PSETAF is of the form $(Arg, \rightsquigarrow, \succ)$, where Arg is a finite set of arguments, the attack relation $\rightsquigarrow$ consists of tuples of the form (S, arg) for nonempty subsets S of Arg and elements arg of Arg, and $\succ$ is an acyclic binary relation on Arg, called the preference relation.

The translation described in [9] uses the ABox as set of arguments and the priority relation as preference relation. The definition of the attack relation is based on the conflict sets. In our adaptation, we use the α-priority relation and the α-conflict sets instead. Following [9], we make the assumption that conflict sets are of cardinality at least two (see [9] for an explanation why this assumption is in principle without loss of generality). This is the case if α is not $\mathcal{T}$-entailed by a single assertion in $\mathcal{A}$.

Definition 7 (Associated PSETAF)**.** *Let $\mathcal{K}_{\succ_\alpha} = (\mathcal{K}, \succ_\alpha)$ be a prioritized KB for $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ and α an assertion that is not $\mathcal{T}$-entailed by any single assertion in $\mathcal{A}$. The* PSETAF associated with $\mathcal{K}_{\succ_a}$ *is the PSETAF $F_{\mathcal{K},\succ_\alpha} := (\mathcal{A}, \rightsquigarrow_\mathcal{K}, \succ_\alpha)$, where $\rightsquigarrow_\mathcal{K} := \{(C \setminus \{\beta\}, \beta) \mid C \in Conf_\alpha(\mathcal{K}), \beta \in C\}$.*

Theorem 35 in [9] states that the Pareto-optimal repairs of an inconsistency are exactly the stable extensions of the PSETAF obtained by the translation in [9]. This result also holds in our setting.

Theorem 1. *Let $\mathcal{K}_{\succ_\alpha} = (\mathcal{K}, \succ_\alpha)$ be a prioritized KB for $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ and α an assertion that is not $\mathcal{T}$-entailed by any single assertion in $\mathcal{A}$. The set $\mathcal{A}' \subseteq \mathcal{A}$ is a Pareto-optimal α-repair of $\mathcal{K}_{\succ_\alpha}$ iff $\mathcal{A}'$ is a stable extension of $F_{\mathcal{K},\succ_\alpha}$.*

Proof sketch. In principle, one could prove this result by redoing the proof of Theorem 35 in [9] in our new setting. However, one can also obtain our Theorem 1 as a corollary to Theorem 35 in [9]. To that purpose, let $\mathcal{K}_{\succ_\alpha}$ with $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ and $\alpha = C(a)$ be a prioritized KB satisfying the prerequisites of the theorem. We modify the $\mathcal{EL}$ TBox $\mathcal{T}$ into an $\mathcal{ELO}^\bot$ TBox $\mathcal{T}'$ by setting $\mathcal{T}' := \mathcal{T} \cup \{\{a\} \sqcap C \sqsubseteq \bot\}$. Let $\mathcal{K}'_{\succ_\alpha} := ((\mathcal{T}', \mathcal{A}), \succ_\alpha)$ be the prioritized $\mathcal{ELO}^\bot$ KB obtained from $\mathcal{K}_{\succ_\alpha}$ by replacing $\mathcal{T}$ with $\mathcal{T}'$. It is easy to see that the following holds for any subset $\mathcal{A}'$ of $\mathcal{A}$:

$$\mathcal{A}' \ \mathcal{T}\text{-entails } C(a) \ \text{ iff } \ \mathcal{A}' \ \mathcal{T}'\text{-entails an inconsistency.}$$

Consequently, an ABox $\mathcal{A}' \subseteq \mathcal{A}$ is an α-conflict w.r.t. $\mathcal{T}$ iff $\mathcal{A}'$ is a conflict w.r.t. $\mathcal{T}'$. In addition, the notion of repair for $\mathcal{K}'_{\succ_\alpha}$ in [9] coincides with our notion of α-repair for $\mathcal{K}_{\succ_\alpha}$, and the same is true for the notion of Pareto-optimal repair and Pareto-optimal α-repair. Applying the translation into a PSETAF of [9] to $\mathcal{K}'_{\succ_\alpha}$ yields exactly the same PSETAF as applying the translation in our Definition 3.2 to $\mathcal{K}_{\succ_\alpha}$. In fact, the employed ABox and the priority relation are the same, and the definition of the attack relation only depends on the conflict sets, which coincide.

It remains to see that Theorem 35 in [9] really applies to $\mathcal{K}'_{\succ_\alpha}$, which is an $\mathcal{ELO}^\perp$ KB rather than a DL-Lite KB. However, if one looks at the proof of Theorem 35 in the extended version [24] of [9], then one sees that no DL-Lite-specific arguments were employed. Thus, this proof also works for $\mathcal{ELO}^\perp$ KBs. □

In the case where the priority relation is the empty relation, Theorem 38 in [9] shows an equivalence between the grounded extension and the intersection of all repairs. In our setting, the analogous result can be formulated as follows and proved by the same argument as employed in the proof of Theorem 1.

Theorem 2. *Let $\mathcal{K}_{\succ_\alpha} = (\mathcal{K}, \succ_\alpha)$ be a prioritized KB for $\mathcal{K} = (\mathcal{T}, \mathcal{A})$, the empty priority relation $\succ_\alpha$, and an assertion α that is not $\mathcal{T}$-entailed by any single assertion in $\mathcal{A}$. Then the unique grounded extension of $F_{\mathcal{K},\succ_\alpha}$ coincides with the intersection of all α-repairs of $\mathcal{K}_{\succ_\alpha}$.*

4. Conclusion

The main purpose of this paper was to point out that, while an inconsistency is the most blatant manifestation of an error in a KB, consistent KBs may still be erroneous. This can be detected when reasoning derives a consequence that, according to the domain knowledge of the user, does not hold in the application domain modelled by the KB. Approaches and results developed to deal with inconsistencies can be adapted to address errors that manifest themselves by the entailment of such an unwanted consequence. Instead of inconsistency-tolerant reasoning one then needs to use error-tolerant reasoning. We have illustrated such adaptations by considering some of the definitions and results of [9], and translating them from the setting of inconsistency-tolerant reasoning for DL-Lite KBs to error-tolerant reasoning for $\mathcal{EL}$ KBs. From a technical point of view, the results shown in the present paper regarding the connection between α-repairs (i.e., repairs that get rid of the unwanted consequence α) and certain extensions of AFs are easy consequences of the corresponding results in [9]. Things get more interesting if, instead of classical repairs (which are subsets of the KB), one considers gentle repairs [16] or optimal repairs [25] (which try to preserve more consequences). While error-tolerant reasoning with respect to optimal repairs has been considered in the setting of $\mathcal{EL}$ [18, 19, 20], it is unclear how to establish a connection to AFs for such more sophisticated notions of repair. We will investigate this problem in our future research.

Acknowledgements

Financial support by the Federal Ministry of Research, Technology and Space of Germany and by Sächsische Staatsministerium für Wissenschaft, Kultur und Tourismus in the program Center of Excellence for AI research “Center for Scalable Data Analytics and Artificial Intelligence Dresden/Leipzig”, project identification number ScaDS.AI is hereby acknowledged.

Generative AI

No use of generative AI tools has been made during the production of this work.

References

[1] N. Rescher, R. Manor, On inference from inconsistent premisses, Theory and Decision 1 (1970) 179–217. doi:10.1007/BF00154005.

[2] J. Chomicki, Consistent query answering: Five easy pieces, in: T. Schwentick, D. Suciu (Eds.), Database Theory - ICDT 2007, 11th International Conference, Barcelona, Spain, January 10-12, 2007, Proceedings, Lecture Notes in Computer Science, Springer, 2007, pp. 1–17. doi:10.1007/11965893_1.

[3] D. Lembo, M. Lenzerini, R. Rosati, M. Ruzzi, D. F. Savo, Inconsistency-tolerant query answering in ontology-based data access, J. Web Semant. 33 (2015) 3–29. doi:10.1016/J.WEBSEM.2015.04.002.

[4] J. Baget, S. Benferhat, Z. Bouraoui, M. Croitoru, M. Mugnier, O. Papini, S. Rocher, K. Tabia, Inconsistency-tolerant query answering: Rationality properties and computational complexity analysis, in: L. Michael, A. C. Kakas (Eds.), Logics in Artificial Intelligence - 15th European Conference, JELIA 2016, Larnaca, Cyprus, November 9-11, 2016, Proceedings, Lecture Notes in Computer Science, 2016, pp. 64–80. doi:10.1007/978-3-319-48758-8_5.

[5] M. Bienvenu, C. Bourgaux, F. Goasdoué, Computing and explaining query answers over inconsistent dl-lite knowledge bases, J. Artif. Intell. Res. 64 (2019) 563–644. doi:10.1613/JAIR.1.11395.

[6] L. Amgoud, S. Kaci, An Argumentation Framework for Merging Conflicting Knowledge Bases: The Prioritized Case, in: L. Godo (Ed.), Symbolic and Quantitative Approaches to Reasoning with Uncertainty, volume 3571, Springer Berlin Heidelberg, 2005, pp. 527–538. URL: http://link.springer.com/10.1007/11518655_45. doi:10.1007/11518655_45.

[7] S. Alejandro Gómez, C. Iván Chesñevar, G. R. Simari, Reasoning with Inconsistent Ontologies Through Argumentation, Applied Artificial Intelligence 24 (2010) 102–148. URL: https://doi.org/10.1080/08839510903448692. doi:10.1080/08839510903448692.

[8] M. Croitoru, S. Vesic, What Can Argumentation Do for Inconsistent Ontology Query Answering?, in: W. Liu, V. S. Subrahmanian, J. Wijsen (Eds.), Scalable Uncertainty Management, volume 8078, Springer Berlin Heidelberg, 2013, pp. 15–29. URL: http://link.springer.com/10.1007/978-3-642-40381-1_2. doi:10.1007/978-3-642-40381-1_2.

[9] M. Bienvenu, C. Bourgaux, Querying and Repairing Inconsistent Prioritized Knowledge Bases: Complexity Analysis and Links with Abstract Argumentation, in: Proceedings of the Seventeenth International Conference on Principles of Knowledge Representation and Reasoning, International Joint Conferences on Artificial Intelligence Organization, 2020, pp. 141–151. doi:10.24963/kr.2020/15.

[10] Y. Mahmood, J. Virtema, T. Barlag, A.-C. N. Ngomo, Computing Repairs Under Functional and Inclusion Dependencies via Argumentation, in: A. Meier, M. Ortiz (Eds.), Foundations of Information and Knowledge Systems, Springer Nature Switzerland, 2024, pp. 23–42. doi:10.1007/978-3-031-56940-1_2.

[11] P. M. Dung, On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games, Artificial Intelligence 77 (1995) 321–357. doi:10.1016/0004-3702(94)00041-X.

[12] P. Baroni, D. Gabbay, M. Giacomin, L. van der Torre (Eds.), Handbook of Formal Argumentation, College Publications, 2018. URL: http://www.collegepublications.co.uk/downloads/handbooks00003.pdf.

[13] G. Flouris, A. Bikakis, A comprehensive study of argumentation frameworks with sets of attacking arguments, Int. J. Approx. Reason. 109 (2019) 55–86. doi:10.1016/J.IJAR.2019.03.006.

[14] B. Suntisrivaraporn, F. Baader, S. Schulz, K. A. Spackman, Replacing SEP-triplets in SNOMED CT using tractable description logic operators, in: R. Bellazzi, A. Abu-Hanna, J. Hunter (Eds.), Artificial Intelligence in Medicine, 11th Conference on Artificial Intelligence in Medicine, AIME 2007, Amsterdam, The Netherlands, July 7-11, 2007, Proceedings, Lecture Notes in Computer Science, Springer, 2007, pp. 287–291. doi:10.1007/978-3-540-73599-1_38.

[15] F. Baader, B. Suntisrivaraporn, Debugging SNOMED CT using axiom pinpointing in the description

logic $\mathcal{EL}^+$, in: R. Cornet, K. A. Spackman (Eds.), Proceedings of the Third International Conference on Knowledge Representation in Medicine, Phoenix, Arizona, USA, May 31st - June 2nd, 2008, CEUR Workshop Proceedings, CEUR-WS.org, 2008. URL: https://ceur-ws.org/Vol-410/Paper01.pdf.

[16] F. Baader, F. Kriegel, A. Nuradiansyah, R. Peñaloza, Making repairs in description logics more gentle, in: M. Thielscher, F. Toni, F. Wolter (Eds.), Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018, AAAI Press, 2018, pp. 319–328. URL: https://aaai.org/papers/36-18056-making-repairs-in-description-logics-more-gentle/.

[17] M. Ludwig, R. Peñaloza, Error-tolerant reasoning in the description logic $\mathcal{EL}$, in: E. Fermé, J. Leite (Eds.), Logics in Artificial Intelligence - 14th European Conference, JELIA 2014, Funchal, Madeira, Portugal, September 24-26, 2014. Proceedings, Lecture Notes in Computer Science, Springer, 2014, pp. 107–121. doi:10.1007/978-3-319-11558-0_8.

[18] F. Baader, F. Kriegel, A. Nuradiansyah, Error-tolerant reasoning in the description logic $\mathcal{EL}$ based on optimal repairs, in: G. Governatori, A. Turhan (Eds.), Rules and Reasoning - 6th International Joint Conference on Rules and Reasoning, RuleML+RR 2022, Berlin, Germany, September 26-28, 2022, Proceedings, Lecture Notes in Computer Science, Springer, 2022, pp. 227–243. doi:10.1007/978-3-031-21541-4_15.

[19] F. Baader, F. Kriegel, A. Nuradiansyah, Treating role assertions as first-class citizens in repair and error-tolerant reasoning, in: J. Hong, M. Lanperne, J. W. Park, T. Cerný, H. Shahriar (Eds.), Proceedings of the 38th ACM/SIGAPP Symposium on Applied Computing, SAC 2023, Tallinn, Estonia, March 27-31, 2023, ACM, 2023, pp. 974–982. doi:10.1145/3555776.3577630.

[20] F. Baader, F. Kriegel, A. Nuradiansyah, Inconsistency- and error-tolerant reasoning w.r.t. optimal repairs of $\mathcal{EL}^\bot$ ontologies, in: A. Meier, M. Ortiz (Eds.), Foundations of Information and Knowledge Systems - 13th International Symposium, FoIKS 2024, Sheffield, UK, April 8-11, 2024, Proceedings, Lecture Notes in Computer Science, Springer, 2024, pp. 3–22. doi:10.1007/978-3-031-56940-1_1.

[21] F. Baader, S. Brandt, C. Lutz, Pushing the EL envelope, in: L. P. Kaelbling, A. Saffiotti (Eds.), Proc.of the 19th Int.Joint Conf.on Artificial Intelligence (IJCAI 2005), Morgan Kaufmann, Los Altos, 2005, pp. 364–369. URL: http://ijcai.org/Proceedings/05/Papers/0372.pdf.

[22] M. Bienvenu, A Short Survey on Inconsistency Handling in Ontology-Mediated Query Answering, KI - Künstliche Intelligenz 34 (2020) 443–451. doi:10.1007/s13218-020-00680-9.

[23] M. Bienvenu, C. Bourgaux, Querying inconsistent prioritized data with ORBITS: Algorithms, implementation, and experiments, in: G. Kern-Isberner, G. Lakemeyer, T. Meyer (Eds.), Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning, KR 2022, 2022, pp. 523–532. URL: https://proceedings.kr.org/2022/54/.

[24] M. Bienvenu, C. Bourgaux, Querying and repairing inconsistent prioritized knowledge bases: Complexity analysis and links with abstract argumentation, CoRR abs/2003.05746 (2020). URL: https://arxiv.org/abs/2003.05746. arXiv:2003.05746.

[25] F. Baader, P. Koopmann, F. Kriegel, Optimal repairs in the description logic $\mathcal{EL}$ revisited, in: Proceedings of the 18th European Conference on Logics in Artificial Intelligence (JELIA 2023), volume 14281 of *Lecture Notes in Computer Science*, Springer, 2023, pp. 11–34. doi:10.1007/978-3-031-43619-2_2.