Towards a Principle-based Framework for Repair Selection in Inconsistent Knowledge Bases

Said Jabbour
CRIL, CNRS UMR 8188,
Université d'Artois
Lens, France
jabbour@cril.fr

Yue Ma
LISN, CNRS UMR 9015,
Université Paris-Saclay
Gif-sur-Yvette, France
ma@lisn.fr

Badran Raddaoui
SAMOVAR, Télécom SudParis,
Institut Polytechnique de Paris,
France
badran.raddaoui@telecom-sudparis.eu

ABSTRACT

This paper investigates a general principle-based framework for retrieving preferred repairs from inconsistent knowledge bases under a broad family of strategies. To begin with, we define a set of principles that ensure rational behaviours of repair selection strategies. Then, we classify the strategies into two basic categories: (i) comparing repairs without requiring formula information; and (ii) comparing repairs based on formula information. Based on this classification, we present several novel repair selection strategies and show that our framework encompasses various existing popular strategies. Through a systematical analysis of these selection strategies using the proposed principles, we conclude that our principles allow for effective discrimination among the strategies. Finally, preliminary experimental results are presented to show the feasibility of our approach.

KEYWORDS

Inconsistency handling; Preferred repairs; Rationality principles

ACM Reference Format:
Said Jabbour, Yue Ma, and Badran Raddaoui. 2024. Towards a Principle-based Framework for Repair Selection in Inconsistent Knowledge Bases. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 9 pages.

1 INTRODUCTION

The study of preferences has a long tradition in various disciplines. It has been recently applied in query answering over databases [30], propositional logic knowledge bases (KBs) [22], description logic KBs [6, 9, 32], and the existential rule language [12]. In this context, existing approaches rely on the concept of repairs, that is, the ⊆-maximal consistent subsets of an inconsistent KB.

The need of retrieving preferred repairs is justified by numerous issues with using all repairs. The first drawback of reasoning with all repairs is due to their usual large numbers in real-world applications [25]. Moreover, repairs are often not equally important in practical applications, for instance, when a data source is more reliable than another or a piece of new information is preferred over earlier ones [11]. We may also prefer one repair over another if the former contains less problematic information [17, 21, 28]. Therefore, in this paper, we turn our attention to how to choose the most preferred repairs among all the potential repairs of an inconsistent KB, while filtering out undesired repairs. Unlike [6], and in line with [7, 9, 12, 32], we aim to guarantee that the retrieved consistent subsets are still ⊆-maximal.

Despite their success, the aforementioned proposals on preferred repairs-based query answering have been generally developed using ad hoc repair selection strategies. A basic strategy is based on the cardinality of repairs. More advanced strategies often use some aggregation functions of formula information (e.g., weight, priority level, inconsistency measure), either provided as system inputs or computed from the given KB. Without being limited to aggregation functions, this paper provides a complementary principle-based framework that allows us to define different repair selection strategies to retrieve preferred repairs from inconsistent KBs. This paper makes the following concrete contributions:

- We propose a set of logical principles that guarantee rational behaviours of repair selection strategies. We also present some (in)compatibility results that ensure which principles can(not) be satisfied together (Section 3).
- We classify the repair selection strategies into two main categories: the strategies that allow for comparing repairs directly without using formula information (Section 4.1); and the selection strategies that exploit formula information to select preferred repairs (Section 4.2). We stress here that our framework is broad enough to encompass several existing popular repair selection strategies and also leads to other novel strategies in particular for the first category.
- We systematically evaluate the different repair selection strategies against the proposed principles (Table 1), showing that these principles allow for an effective discrimination among different repair selection strategies.
  We also show that satisfying the same principles does not necessary guarantee agreement on the preferred repair sets. This confirms a similar finding regarding argument ranking within abstract argumentation frameworks [2].

2 FORMAL PRELIMINARIES

Our framework is not restricted to a particular logic language. For a convenient illustration of the notions presented in this paper, we use propositional logic (PL) and description logic DL-Lite.

Propositional Logic (PL) We consider a propositional language over a finite set of propositional variables V, the logical connectives,

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand.

and the truth constants *true* ($\top$) and *false* ($\bot$). We use Greek letters α, β, etc. to denote arbitrary formulas. We further denote by $\vdash$ the classical inference relation in PL. A PL KB $\mathcal{K}$ is a non-empty finite set of propositional formulas. $\mathcal{K}$ is said to be *inconsistent* if $\exists\alpha$ s.t. $\mathcal{K} \vdash \alpha$ and $\mathcal{K} \vdash \neg\alpha$. Let $\Omega = \{w_1, w_2, \ldots\}$ denotes the set of classical interpretations associated with the language, where w_i is a mapping from V to $\{0, 1\}$. An interpretation $w \in \Omega$ is a *model* of a formula α, denoted $w \models \alpha$, if α is *true* in w in the usual way.

DL-Lite Logic Concepts and roles of DL-Lite are formed by the following syntax [13]: $B ::= A \mid \exists R, \quad R ::= P \mid P^-, \quad C ::= B \mid \neg B$. A and P denote an *atomic concept* and an *atomic role*, respectively; B denotes a *basic concept* (i.e., a concept of the form A, $\exists R$); R denotes a *basic role* (a role of the form P, P^-), where P^- denotes the inverse of the atomic role P; C denotes a *general concept* (a concept of the form B, $\neg B$). A DL-Lite TBox (denoted $\mathcal{T}$) is a set of inclusion axioms of the form $B \sqsubseteq C$. A DL-Lite ABox (denoted $\mathcal{A}$) is a set of membership assertions on atomic concepts and atomic roles: $A(a)$, $P(a, b)$ called ABox axioms, where a, b are individuals, saying that a is an instances of A, and a has P relation with b. The semantics of DL-Lite is given by an interpretation $I = (\Delta^I, \cdot^I)$, consisting of a non-empty *interpretation domain* Δ^I and an *interpretation function* $\cdot^I$ that assigns to each concept C a subset C^I of Δ^I, and to each role R a binary relation R^I over $\Delta^I \times \Delta^I$. The extension of $\cdot^I$ to general concepts follows the usual way. An interpretation satisfies a DL-Lite KB $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ (i.e., a *model* of $\mathcal{K}$) iff it satisfies each axiom in both *ABox* and *TBox*. A KB is satisfiable if it has at least one model. A KB $\mathcal{K}$ logically implies an assertion α, written $\mathcal{K} \models \alpha$, if all models of $\mathcal{K}$ are also models of α.

Formula and Knowledge Base In this paper, we use *formulas* to refer to both DL-Lite axioms and PL formulas when no distinction of language is necessary. Let $\mathcal{L}$ be DL-Lite or PL and $\mathsf{Form}(\mathcal{L})$ be the infinite set of formulas over $\mathcal{L}$. A knowledge base is a finite subset of $\mathsf{Form}(\mathcal{L})$. We write $\mathbb{K}$ for the set of all knowledge bases in $\mathcal{L}$. Given $\{\alpha, \beta\} \subseteq \mathsf{Form}(\mathcal{L})$, α and β are *logically equivalent*, denoted $\alpha \equiv \beta$, if $\{\alpha\} \models \beta$ and $\{\beta\} \models \alpha$. Now, let us recall the equivalence relation between sets of formulas [3].

Definition 2.1. Let $\mathcal{K}, \mathcal{K}' \in \mathbb{K}$. Then, $\mathcal{K}$ and $\mathcal{K}'$ are **logically equivalent**, denoted $\mathcal{K} \cong \mathcal{K}'$, iff $\forall\alpha \in \mathcal{K}$, there exists $\beta \in \mathcal{K}'$ s.t. $\alpha \equiv \beta$, and $\forall\alpha' \in \mathcal{K}'$, there exists $\beta' \in \mathcal{K}$ s.t. $\alpha' \equiv \beta'$.

Repairs and Conflicts Handling inconsistencies is a vital aspect in many intelligent systems, relying on two main notions, namely the *repairs* and the *conflicts*. Formally, they are defined as follows.

Definition 2.2. Let $\mathcal{K}$ be a PL KB, and $\mathcal{K}'$ be a subset of $\mathcal{K}$. We say $\mathcal{K}'$ a **repair** (resp. **conflict**) of $\mathcal{K}$ if $\mathcal{K}'$ is a $\subseteq$-maximal consistent (resp. $\subseteq$-minimal inconsistent) subset of $\mathcal{K}$. For a DL KB $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$, an ABox **repair** is a $\subseteq$-maximal subset $\mathcal{A}'$ of $\mathcal{A}$ such that $\mathcal{T} \cup \mathcal{A}'$ is consistent. An ABox **conflict** is a $\subseteq$-minimal subset $\mathcal{A}''$ of $\mathcal{A}$ s.t. $\langle \mathcal{T}, \mathcal{A}'' \rangle \models \bot$. We denote by $\mathcal{R}(\mathcal{K})$ and $C(\mathcal{K})$ the set of all possible repairs and conflicts of $\mathcal{K}$, respectively.

In this way, if $\mathcal{K}_1 \cong \mathcal{K}_2$, then $|\mathcal{R}(\mathcal{K}_1)| = |\mathcal{R}(\mathcal{K}_2)|$. In what follows, we define $\mathcal{R}(\alpha, \mathcal{K}) = \{R \in \mathcal{R}(\mathcal{K}) \mid \alpha \in R\}$ and $C(\alpha, \mathcal{K}) = \{C \in C(\mathcal{K}) \mid \alpha \in C\}$. A formula $\alpha \in \mathcal{K}$ is **free** iff $C(\alpha, \mathcal{K}) = \emptyset$. We denote by $\mathsf{Free}(\mathcal{K})$ the set of free formulas in $\mathcal{K}$. Let us note that a single formula can be a conflict (named a *self-contradictory* formula), e.g., $A(a)$ w.r.t. $\mathcal{T} = \{\langle A \sqsubseteq B, A \sqsubseteq \neg B \rangle\}$. If each formula of $\mathcal{K}$ gives a conflict, then there will be no repair in $\mathcal{K}$, i.e., $\mathcal{R}(\mathcal{K}) = \emptyset$. Removing such formulas by existing reasoning algorithms makes the KB they belong to empty and makes it no sense of selecting repairs [7]. From now on, without loss of generality and for the sake of simplicity, we shall assume that there is no self-contradictory formulas nor tautologies.

Order and Preorder Let X be a finite set of objects. A binary relation is called an *order* if it is reflexive, transitive and total. A *preorder* is a reflexive and transitive relation. If $\geq$ is an order on X, the induced *strict* order $>$ on X is defined by $x > y$ if $x \geq y$ and $y \not\geq x$, where x, y are arbitrary elements of X. A binary relation is said to be *acyclic* if it does not contains cycles. We further say that $>^c$ is a *completion* of $>$ if $>^c$ is a total strict order and for any $x > y$, we have $x >^c y$.

3 PRINCIPLE-BASED FRAMEWORK FOR REPAIR SELECTION

Several inconsistency-tolerant semantics have been recently extended to query inconsistent knowledge and databases under preferred repairs, e.g., [6, 9, 12, 22, 30, 32]. Despite their potentially fruitful success, these proposals have been generally developed using basic strategies to retrieve the most preferred repairs. In fact, the strategies are often some aggregation functions of formula information (e.g., weight, priority level), either provided as a system parameter or computed from the input KB. That is why the main goal in this paper is to define a complementary principle-based framework to filter the slate of preferred repairs that can be used to draw meaningful results from inconsistent KBs when consistency cannot be restored. Such framework encompasses most of the existing methods for selecting preferred repairs, which have been previously used for both Database and Logic areas. Note here that there is no single best strategy to express preferences among sets [10]. In what follows, we introduce a broad family of selection repair strategies which allows to identify the most appropriate strategy according to the application context. One of our motivations for this work is based on the following observation:

Proposition 3.1. *Let $\mathcal{K} \in \mathbb{K}$. $\mathcal{K}$ is inconsistent iff $|\mathcal{R}(\mathcal{K})| \geq 2$.*

That is, any inconsistent KB has multiple repairs, which makes the task of comparing these repairs not obvious. In what follows, the repairs of a given KB will be ordered according to some criteria to determine only the most desired ones. More formally, our framework returns the set of *preferred repairs* among the (large) candidate set of repairs by taking as input three basic elements:

- a knowledge base $\mathcal{K} \in \mathbb{K}$,
- a *formula characterisation ch* over formulas in $\mathcal{L}$: *ch* can be a binary function or relation over $\mathsf{Form}(\mathcal{L})$, e.g., a weight, a distance, or a priority relation between two formulas. Though *ch* could be arbitrary, the following property is natural for *ch* to satisfy: If $\alpha \equiv \alpha'$ and $\beta \equiv \beta'$, then $ch(\alpha, \beta) = ch(\alpha', \beta')$.
- a *repair comparison strategy*, written $\succcurlyeq_s \subseteq \mathcal{R}(\mathcal{K}) \times \mathcal{R}(\mathcal{K})$, to compare the repairs of $\mathcal{K}$.

Note that the selection repair strategies defined in Section 4.1 are independent of any specific *ch*, while the strategies based on categories of *ch* are discussed in Section 4.2.

Definition 3.2. Let $\mathcal{K} \in \mathbb{K}$ and $R, R' \in \mathcal{R}(\mathcal{K})$. A **repair comparison strategy** $\succcurlyeq_s$ is an acyclic preference relation over $\mathcal{R}(\mathcal{K})$ with $R \succcurlyeq_s R'$, meaning that R is preferred over R'. Given a relation ch, a **repair selection function** is a mapping $\mathcal{F} : \mathcal{K} \times ch \times \succcurlyeq_s \rightarrow \Xi$ s.t. $\Xi \in 2^{\mathcal{R}(\mathcal{K})}$.

A classical way to select repairs is to pick the preferred repairs according to the repair comparison strategy $\succcurlyeq_s$.

Definition 3.3. Given $\mathcal{K} \in \mathbb{K}$, an **optimal repair selection function** is a repair selection function $\mathcal{F}$ such that there is no $R \in \mathcal{F}(\mathcal{K}, ch, \succcurlyeq_s)$ and $R' \in (\mathcal{R}(\mathcal{K}) \setminus \mathcal{F}(\mathcal{K}, ch, \succcurlyeq_s))$ such that $R' \succ_s R$.

Note that the acyclicness of $\succcurlyeq_s$ guarantees the existence of a set of selected repairs by $\mathcal{F}$. If the repair comparison strategy $\succcurlyeq_s$ is a strict total relation (i.e., transitive and anti-reflexive), besides the optimal repair selection, one can also define a selection function that outputs repairs ranking on the Top-k positions according to $\succcurlyeq_s$. Due to the space constraint, we focus in this paper on optimal repair selection function. When the comparison strategy is a strict partial order, the optimal repair selection returns the "skyline".

Rationality Principles

Definition 3.3 is general enough regarding the repair comparison strategy $\succcurlyeq_s$ and the formula characterisation ch. To restrict the possible candidates, we first establish a set of desiderata that a suitable optimal repair selection function $\mathcal{F}$ should satisfy. Such formal principles are important for defining, characterizing, and comparing selection functions. Below, we introduce the set of principles[1].
Abstraction: For any $\mathcal{K} \in \mathbb{K}$ and any isomorphism[2] γ, we have that $\forall R \in \mathcal{R}(\mathcal{K})$, $R \in \mathcal{F}(\mathcal{K}, ch, \succcurlyeq_s)$ iff $\gamma(R) \in \mathcal{F}(\gamma(\mathcal{K}), ch, \succcurlyeq_s)$.
Equivalence Invariance: If $\mathcal{K}_1 \cong \mathcal{K}_2$, then for each $R \in \mathcal{F}(\mathcal{K}_1, ch, \succcurlyeq_s)$, there exists $R' \in \mathcal{F}(\mathcal{K}_2, ch, \succcurlyeq_s)$ s.t. $R \cong R'$.
Coverness: For any KB $\mathcal{K}$, $\bigcup_{R \in \mathcal{F}(\mathcal{K}, ch, \succcurlyeq_s)} R = \mathcal{K}$.
Non-Emptiness: $\mathcal{F}(\mathcal{K}, ch, \succcurlyeq_s) \neq \emptyset$, if $\mathcal{R}(\mathcal{K}) \neq \emptyset$.
Monotonicity: If ch_1, ch_2 are two relations s.t. $ch_2 \subseteq ch_1$, then $\mathcal{F}(\mathcal{K}, ch_1, \succcurlyeq_s) \subseteq \mathcal{F}(\mathcal{K}, ch_2, \succcurlyeq_s)$.
Non-Discrimination: $\mathcal{F}(\mathcal{K}, ch, \emptyset) = \mathcal{R}(\mathcal{K})$.
Reducibility: If $\mathcal{K} \vdash \bot$, then $\mathcal{F}(\mathcal{K}, ch, \succcurlyeq_s) \subset \mathcal{R}(\mathcal{K})$.
Consistency: If $\mathcal{K} \nvdash \bot$, then $\mathcal{F}(\mathcal{K}, ch, \succcurlyeq_s) = \{\mathcal{K}\}$.
Improvement: Given $\mathcal{K}, \mathcal{K}' \in \mathbb{K}$, if $\mathcal{K} \subseteq \mathcal{K}'$, then for any $R \in \mathcal{F}(\mathcal{K}, ch, \succcurlyeq_s)$ and $R' \in \mathcal{F}(\mathcal{K}', ch, \succcurlyeq_s)$, $R \not\succ_s R'$.
Stability: Given $\mathcal{K}, \mathcal{K}' \in \mathbb{K}$, if $\mathcal{K} \subseteq \mathcal{K}'$, for any $R \in \mathcal{F}(\mathcal{K}, ch, \succcurlyeq_s)$, there exists $R' \in \mathcal{F}(\mathcal{K}', ch, \succcurlyeq_s)$ such that $R \subseteq R'$.
Uniqueness: $|\mathcal{F}(\mathcal{K}, ch, \succcurlyeq_s)| = 1$.
Non-Triviality: There is $\mathcal{K} \in \mathbb{K}$ such that $\mathcal{F}(\mathcal{K}, ch, \succcurlyeq_s) \subset \mathcal{R}(\mathcal{K})$.

The intuition about the principles is as follows: The *Abstraction* principle states that the selection of preferred repairs should be independent from the variables' names. The *Equivalence Invariance* concerns the main intuition behind the logical equivalence between two sets of formulas. It ensures that two equivalent KBs exhibit the same behaviour according to the function $\mathcal{F}$. The third principle, called *Coverness*, requires that a KB should be covered by its preferred repairs, i.e., each formula in $\mathcal{K}$ must belong to at least one preferred repair. The *Non-Emptiness* principle says that an optimal repair selection strategy $\mathcal{F}$ must return at least one preferred repair. *Monotonicity* implies that looking at richer preferences among formulas can only narrow down the set of preferred repairs. The *Non-Discrimination* principle states that the preferred repairs are the set of all repairs when no repair comparison strategy is expressed. *Reducibility* claims that there always exists a manner to remove certain repairs from being the selected preferred ones. *Consistency* states that if a KB is consistent, then the only preferred repair is the base itself. The *Improvement* principle says that the expansion of a KB can only result in better preferred repairs according to $\succcurlyeq_s$. The *Stability* principle states that the expansion of a KB $\mathcal{K}$ preserves all the formulas involved in the preferred repairs of $\mathcal{K}$. *Uniqueness* says that an optimal repair selection strategy should return one unique preferred repair. Finally, the *Non-Triviality* principle is to avoid repair selection functions that do not discriminate between repairs and return all repairs for any KB $\mathcal{K}$.

Proposition 3.4. *Given an optimal repair selection function $\mathcal{F}$, if $\mathcal{F}$ satisfies Uniqueness, then it satisfies Reducibility. If $\mathcal{F}$ satisfies Reducibility, it satisfies Non-Triviality. But, the converses are false.*

The following result shows that some principles are not necessarily compatible with one another, that is, they cannot hold together for the same optimal repair selection function.

Proposition 3.5. *Let $\mathcal{K} \in \mathbb{K}$ s.t. $\mathcal{K} \vdash \bot$. There exists no optimal repair selection function $\mathcal{F}$ which satisfies both Coverness and Uniqueness. Moreover, there exists no optimal repair selection function $\mathcal{F}$ which satisfies both Non-Discrimination and Uniqueness.*

Below, we present a sufficient condition to satisfy the *Improvement* principle via the notion of *Subset-Dependence*:

Definition 3.6. We say that a binary relation rel satisfies the *Subset-Dependence* property if $rel(X, Y)$ holds for any sets $Y \subseteq X$.

Proposition 3.7. *If a repair comparison strategy $\succcurlyeq_s$ is transitive and satisfies the* Subset-Dependence *property, then the optimal repair selection function $\mathcal{F}$ satisfies the Improvement principle.*

Hopefully, there are nine compatible principles, i.e., they can be satisfied all together by an optimal repair selection function.

Proposition 3.8. *The properties Abstraction, Equivalence Invariance, Non-Emptiness, Monotonicity, Non-Discrimination, Consistency, Improvement, Stability and Non-Triviality are compatible.*

Note that other combinations of compatible principles can be found in Table 1 that enumerates the principles satisfied by the optimal repair selection functions discussed below.

4 SELECTION OF PREFERRED REPAIRS

This section presents various optimal repair selection functions to identify the most preferred repairs. Roughly speaking, these functions can be classified into two main categories: the first category is intended for the setting where no formula information is available; whilst the selection functions from the second category takes formulas information into account during the selection process.

By abusing the notation a little bit without confusion, in the sequel, we also use $\succcurlyeq_s$ when we refer to $\mathcal{F}(\mathcal{K}, ch, \succcurlyeq_s)$. In turn, we

[1] The three principles *Non-Emptiness*, *Monotonicity*, and *Non-Discrimination* have been proposed by [30] for preferred consistent query answers in relational databases.
[2] The isomorphism γ renames all the variables (resp. concept, role names, and instances) for PL (resp. DL). We apply γ to formulas and knowledge bases.

simply call $\succcurlyeq_s$ a *repair selection strategy*. Sometimes, we use $\succcurlyeq_s^x$ to indicate a specific selection strategy from the family s with a sub-type x (e.g. see Definition 4.5).

4.1 *ch*-Free Selection Strategies

This subsection introduces our strategies for retrieving preferred repairs without using the relation *ch* (i.e., without any extra information about the formulas of the KB). Obviously, in that case, the *Monotonicity* principle is always satisfied.

4.1.1 Cardinality-based strategy. A straightforward way for filtering preferred repairs involves employing a cardinality-based criterion [9] that can be formalized in our notation as follows. Given two repairs $R, R' \in \mathcal{R}(\mathcal{K})$, $R \succcurlyeq_{\text{card}} R'$ iff. $|R| \geq |R'|$. The key idea underlying the cardinality-based method is to retain as many information as possible.

Proposition 4.1. *$\succcurlyeq_{\text{card}}$ satisfies Abstraction, Non-Emptiness, Monotonicity, Non-Discrimination, Consistency, Improvement, and Non-Triviality principles. However, $\succcurlyeq_{\text{card}}$ does not satisfy Equivalence Invariance, Coverness, Reducibility, Stability or Uniqueness. For DL-Lite, $\succcurlyeq_{\text{card}}$ satisfies the Equivalence Invariance principle.*

Remark. Given $\mathcal{K}, \mathcal{K}' \in \mathbb{K}$, $R \in \mathcal{R}(\mathcal{K})$ and $R' \in \mathcal{R}(\mathcal{K}')$ satisfying $\mathcal{K} \cong \mathcal{K}'$ and $R \cong R'$, it does not always hold that $|R| = |R'|$. This is because one formula in $\mathcal{K}_1$ can have several logically equivalent formulas in $\mathcal{K}_2$. Hence, $\succcurlyeq_{\text{card}}$ does not satisfy *Equivalence Invariance.*

The cardinality-based strategy $\succcurlyeq_{\text{card}}$ treats all repairs independently. A problem with this strategy is that no account is taken of the possible interaction between repairs of a given KB.

4.1.2 Compatibility-based strategy. This second strategy, which we call *compatibility-based strategy*, compares all pairs of repairs based on the next criterion: we prefer the repairs having more compatibility with other repairs, i.e., opposed by less repairs. Formally,

Definition 4.2. Let $\mathcal{K} \in \mathbb{K}$ and $R \in \mathcal{R}(\mathcal{K})$. Then, the ***compatible set*** of R w.r.t. $\mathcal{R}(\mathcal{K})$, written $\text{comp}(R, \mathcal{K})$, is defined as $\text{comp}(R, \mathcal{K}) = \{R' \in \mathcal{R}(\mathcal{K}) \mid (R \setminus R') \cup (R' \setminus R) \nvdash \bot\}$.

Since a repair is a maximal consistent subset, adding any extra formula will lead to inconsistency directly. However, it is still relevant to consider the elements that distinguish the two repairs without common parts (*i.e.*, $(R \setminus R') \cup (R' \setminus R)$) as claimed in [10].

Example 4.3. Let $\mathcal{K} = \{a, \neg a, \neg b, \neg a \vee b\}$. $\mathcal{K}$ has three repairs $R_1 = \{\neg a, \neg b, \neg a \vee b\}$, $R_2 = \{a, \neg a \vee b\}$, and $R_3 = \{a, \neg b\}$. Then, $\text{comp}(R_1, \mathcal{K}) = \{R_1\}$ and $\text{comp}(R_2, \mathcal{K}) = \text{comp}(R_3, \mathcal{K}) = \{R_2, R_3\}$.

It is easy to see from Definition 4.2 that each repair belongs to the compatible set of itself. A repair R is contained in the compatible set of a repair R' iff R' is also part of the compatible set of R.

Corollary 4.4. *Let $\mathcal{K} \in \mathbb{K}$. Given two repairs $R, R' \in \mathcal{R}(\mathcal{K})$, it holds that $R \in \text{comp}(R, \mathcal{K})$. Also, if $R' \in \text{comp}(R, \mathcal{K})$, then $R \in \text{comp}(R', \mathcal{K})$.*

Having defined the compatible sets of repairs, we are now ready to introduce our compatibility-based strategy:

Definition 4.5. Let $\mathcal{K} \in \mathbb{K}$. For two repairs $R, R' \in \mathcal{R}(\mathcal{K})$, we say that R is preferred to R' in terms of cardinality (resp. set-inclusion), written $R \succcurlyeq^{\#}_{\text{comp}} R'$ (resp. $R \succcurlyeq^{\supseteq}_{\text{comp}} R'$), iff: $|\text{comp}(R, \mathcal{K})| \geq |\text{comp}(R', \mathcal{K})|$ (resp. $\text{comp}(R, \mathcal{K}) \supseteq \text{comp}(R', \mathcal{K})$).

This definition can be generalized to compare two repairs $R \in \mathcal{R}(\mathcal{K})$ and $R' \in \mathcal{R}(\mathcal{K}')$ of two KBs $\mathcal{K}$ and $\mathcal{K}'$, which is necessary to study the *Improvement* principle, as follows: $R \succcurlyeq^{\#}_{\text{comp}} R'$ (resp. $R \succcurlyeq^{\supseteq}_{\text{comp}} R'$), if $|\text{comp}(R, \mathcal{K})| \geq |\text{comp}(R', \mathcal{K}')|$ (resp. $\text{comp}(R, \mathcal{K}) \supseteq \text{comp}(R', \mathcal{K}')$).

The following two examples show that the compatibility-based strategies are different from $\succcurlyeq_{\text{card}}$ and $\mathcal{R}(\mathcal{K})$.

Example 4.6 (Example 4.3 contd.). We have $\mathcal{F}(\mathcal{K}, \emptyset, \succcurlyeq_{\text{card}}) = \{R_1\}$, $\mathcal{F}(\mathcal{K}, \emptyset, \succcurlyeq^{\#}_{\text{comp}}) = \{R_2, R_3\}$, and $\mathcal{F}(\mathcal{K}, \emptyset, \succcurlyeq^{\supseteq}_{\text{comp}}) = \{R_1, R_2, R_3\}$.

Example 4.7. Let $\mathcal{K} = \{\alpha_1, \alpha_2, \alpha_3, \alpha_4, \alpha_5, \alpha_6\}$ with $\alpha_1 = a \wedge d$, $\alpha_2 = \neg a \wedge e$, $\alpha_3 = \neg a \vee f$, $\alpha_4 = \neg f \vee \neg d$, $\alpha_5 = \neg e \vee c$, $\alpha_6 = \neg b \wedge \neg c$. $\mathcal{K}$ has three conflicts: $\{\alpha_1, \alpha_3, \alpha_4\}$, $\{\alpha_1, \alpha_2\}$, $\{\alpha_2, \alpha_5, \alpha_6\}$, and five repairs: $R_1 = \{\alpha_1, \alpha_3, \alpha_5, \alpha_6\}$, $R_2 = \{\alpha_1, \alpha_4, \alpha_5, \alpha_6\}$, $R_3 = \{\alpha_2, \alpha_3, \alpha_4, \alpha_5\}$, $R_4 = \{\alpha_2, \alpha_3, \alpha_4, \alpha_6\}$, $R_5 = \{\alpha_3, \alpha_4, \alpha_5, \alpha_6\}$. Then, $\text{comp}(R_1, \mathcal{K}) = \text{comp}(R_2, \mathcal{K}) = \{R_1, R_2, R_5\}$, $\text{comp}(R_3, \mathcal{K}) = \text{comp}(R_4, \mathcal{K}) = \{R_3, R_4, R_5\}$, and $\text{comp}(R_5, \mathcal{K}) = \{R_1, R_2, R_3, R_4, R_5\}$. So $\mathcal{F}(\mathcal{K}, \emptyset, \succcurlyeq^{\supseteq}_{\text{comp}}) = \{R_5\} \subset \mathcal{R}(\mathcal{K})$.

The compatibility-based strategy $\succcurlyeq^x_{\text{comp}}$ $(x \in \{\#, \supseteq\})$ extends $\succcurlyeq_{\text{card}}$ according to the following result:

Proposition 4.8. *Let $\mathcal{K} \in \mathbb{K}$. If $R \cap R' = \emptyset$ for any $R, R' \in \mathcal{R}(\mathcal{K})$, then $\mathcal{F}(\mathcal{K}, \emptyset, \succcurlyeq^{\#}_{\text{comp}}) = \mathcal{F}(\mathcal{K}, \emptyset, \succcurlyeq_{\text{card}})$ and $\mathcal{F}(\mathcal{K}, \emptyset, \succcurlyeq^{\supseteq}_{\text{comp}}) = \mathcal{R}(\mathcal{K})$.*

That is, if all repairs of $\mathcal{K}$ are pairwise disjoint, then $\text{comp}(R, \mathcal{K}) = \{R\}$, thus $\succcurlyeq_{\text{card}}$ and $\succcurlyeq^{\#}_{\text{comp}}$ coincide, and $\succcurlyeq^{\supseteq}_{\text{comp}}$ can be reduced to no selection among repairs.

Proposition 4.9. *$\succcurlyeq^x_{\text{comp}}$ for $x \in \{\#, \supseteq\}$ satisfies Abstraction, Equivalence Invariance, Non-Emptiness, Non-Discrimination, Monotonicity, and Consistency properties. However, $\succcurlyeq^x_{\text{comp}}$ does not satisfy Coverness, Reducibility, Stability, or Uniqueness. The Improvement and Stability holds for DL-Lite, but Non-Triviality is not satisfied by DL-Lite.*

Examples 4.6 and 4.7 are counter-examples for *Coverness* and *Uniqueness.* The counter-examples for *Reductibility, Improvement* and *Stability* are given below in Examples 4.10 and 4.11.

Example 4.10. Let $\mathcal{K}_0 = \{a, \neg b, \neg a \vee b\}$. Then, $\mathcal{K}_0$ has three repairs $R_1 = \{\neg b, \neg a \vee b\}$, $R_2 = \{a, \neg a \vee b\}$, and $R_3 = \{a, \neg b\}$. Thus, we have that the three repairs are pairwise compatible, i.e., $\text{comp}(R_1, \mathcal{K}) = \text{comp}(R_2, \mathcal{K}) = \text{comp}(R_3, \mathcal{K}) = \{R_1, R_2, R_3\}$. So, $\mathcal{F}(\mathcal{K}, \emptyset, \succcurlyeq^{\supseteq}_{\text{comp}}) = \mathcal{F}(\mathcal{K}, \emptyset, \succcurlyeq^{\#}_{\text{comp}}) = \{R_1, R_2, R_3\}$. Consequently, $\succcurlyeq^{\supseteq}_{\text{comp}}$ and $\succcurlyeq^{\#}_{\text{comp}}$ do not satisfy Reducibility.

Example 4.11 (Example 4.10 contd.). Let $\mathcal{K}' = \mathcal{K}_0 \cup \{a \wedge b\}$. $\mathcal{K}'$ has three repairs: $R'_1 = R_1$, $R'_2 = R_2 \cup \{a \wedge b\}$, and $R'_3 = R_3$. Moreover, $\text{comp}(R'_1, \mathcal{K}') = \{R_1, R_3\}$, $\text{comp}(R'_2, \mathcal{K}') = \{R'_2\}$, and $\text{comp}(R'_3, \mathcal{K}') = \{R_1, R_3\}$. So, $\mathcal{F}(\mathcal{K}', \emptyset, \succcurlyeq^{\supseteq}_{\text{comp}}) = \{R_1, R'_2, R_3\}$ and $\mathcal{F}(\mathcal{K}', \emptyset, \succcurlyeq^{\#}_{\text{comp}}) = \{R'_1, R'_3\}$. Recall that $\mathcal{F}(\mathcal{K}, \emptyset, \succcurlyeq^{\supseteq}_{\text{comp}}) = \{R_1, R_2, R_3\}$. Note that $\text{comp}(R_3, \mathcal{K}) \supset \text{comp}(R'_3, \mathcal{K}')$. By definition, $R_3 \succcurlyeq^x_{\text{comp}} R'_3$ $(x \in \{\supseteq, \#\})$, which violates the Improvement. $R'_2 \notin \mathcal{F}(\mathcal{K}', \emptyset, \succcurlyeq^{\#}_{\text{comp}})$, so the Stability is unsatisfied.

Unfortunately, the compatibility-based strategies are not specifically geared toward DL-Lite. In fact, the notion of compatible set of a repair is trivial because it becomes a singleton and only contains the repair itself. This result also explains the violation of Non-Triviality and satisfaction of Stability of $\succcurlyeq^{x}_{\text{comp}}$ ($x \in \{\#, \subseteq\}$) for DL-Lite.

Proposition 4.12. *Given a KB $\mathcal{K} \in \mathbb{K}$, if $\forall C \in C(\mathcal{K})$, $|C| = 2$, then* $\text{comp}(R, \mathcal{K}) = \{R\}$*, $\forall R \in \mathcal{R}(\mathcal{K})$.*

4.1.3 Cover-based strategy. Our third strategy for selecting the most relevant repairs, when no preference among formulas is expressed, is coined the *cover-based strategy*. Generally speaking, one could only retrieve a set of repairs involving the most formulas from the original KB. Then, this strategy aims to selecting among the repairs the ones covering the KB. For that purpose, let us first define the concept of coverness in KBs. Given a KB $\mathcal{K}$, $\Gamma = \{\mathcal{K}_1, \ldots, \mathcal{K}_n \mid \mathcal{K}_i \subseteq \mathcal{K}, 1 \leq i \leq n\}$ is a *cover* of $\mathcal{K}$ iff $\bigcup_{1 \leq i \leq n} \mathcal{K}_i = \mathcal{K}$. A cover Γ of $\mathcal{K}$ is minimal if there exists no other cover Γ' of $\mathcal{K}$ s.t. $\Gamma' \subset \Gamma$. We call Γ a **minimal repair cover** of $\mathcal{K}$, if $\mathcal{K}_i \in \mathcal{R}(\mathcal{K})$ for $1 \leq i \leq n$. We write $\text{cover}(\mathcal{K})$ for the set of minimal repair covers of $\mathcal{K}$.

Example 4.13. Let $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ with $\mathcal{T} = \{A_i \sqsubseteq \neg A'_j, A'_i \sqsubseteq \neg A'_j \mid i \neq j, 1 \leq i, j \leq 3\}$ and $\mathcal{A} = \{A_i(a), A'_i(a) \mid 1 \leq i \leq 3\}$. $\mathcal{K}$ has four repairs: $R_i = \{A_i(a), A'_i(a)\}$ $(1 \leq i \leq 3)$ and $R_4 = \{A_1(a), A_2(a), A_3(a)\}$. Note that R_4 has the biggest cardinality, but it is not in any minimal repair cover.

We are in a position to define our cover-based strategy.

Definition 4.14. Let $\mathcal{K} \in \mathbb{K}$. For two repairs $R, R' \in \mathcal{R}(\mathcal{K})$, we say that $R \succcurlyeq_{\text{cover}} R'$ if there is $\Gamma \in \text{cover}(\mathcal{K})$ such that $R \in \Gamma$.

This definition implies that the set of the preferred repairs will be the union of minimal covers, i.e., $\mathcal{F}(\mathcal{K}, \emptyset, \succcurlyeq_{\text{cover}}) = \bigcup_{\Gamma \in \text{cover}(\mathcal{K})} \Gamma$.

Proposition 4.15. *$\succcurlyeq_{\text{cover}}$ satisfies Abstraction, Equivalence Invariance, Coverness, Non-Emptiness, Monotonicity, Non-Discrimination, Consistency, and Non-Triviality principles. However, $\succcurlyeq_{\text{cover}}$ does not satisfy Reducibility, Improvement, Stability, or Uniqueness.*

Example 4.16 (Example 4.10 contd.). $\mathcal{K}_0$ has three minimal repair covers, namely $\Gamma_1 = \{R_1, R_2\}$, $\Gamma_2 = \{R_1, R_3\}$, and $\Gamma_3 = \{R_2, R_3\}$. Then, $\mathcal{F}(\mathcal{K}, \emptyset, \succcurlyeq_{\text{cover}}) = \Gamma_1 \cup \Gamma_2 \cup \Gamma_3 \{R_1, R_2, R_3\} = \mathcal{R}(\mathcal{K})$. Consequently, $\succcurlyeq_{\text{cover}}$ does not satisfy Reducibility.

Example 4.17 (Example 4.7 contd.). $\mathcal{K}$ has four minimal repair covers: $\{R_1, R_3\}, \{R_2, R_3\}, \{R_1, R_4\}, \{R_2, R_4\}$. Therefore, $\mathcal{F}(\mathcal{K}, \emptyset, \succcurlyeq_{\text{cover}}) = \{R_1, R_2, R_3, R_4\}$. Let $\mathcal{K}' = \mathcal{K} \cup \{b, e \wedge \neg c\}$. $\mathcal{K}'$ has 9 repairs. Three are the same as that of $\mathcal{K}$: $R'_1 = R_1, R'_2 = R_2, R'_5 = R_5$; two are super-sets: $R'_3 = R_3 \cup \{N_7\}$, $R'_4 = R_4 \cup \{N_8\}$; and four new repairs $R'_6 = \{\alpha_1, \alpha_3, \alpha_5, \alpha_7\}, R'_7 = \{\alpha_1, \alpha_4, \alpha_5, \alpha_7\}, R'_8 = \{\alpha_1, \alpha_3, \alpha_6, \alpha_8\}$, and $R'_9 = \{\alpha_1, \alpha_4, \alpha_6, \alpha_8\}$. Then, the minimal covers for $\mathcal{K}'$ are $\{R'_3, R'_8\}$, $\{R'_4, R'_6\}$, $\{R'_4, R'_7\}$, $\{R'_3, R'_9\}$, and $\{R'_4, R'_7\}$, thus $\mathcal{F}(\mathcal{K}', \emptyset, \succcurlyeq_{\text{cover}}) = \{R'_3, R'_4, R'_8, R'_6, R'_7, R'_9\}$. Hence, $\succcurlyeq_{\text{cover}}$ does not satisfy Stability, but it satisfies Non-Triviality.

Similarly, we can construct an example showing that $\succcurlyeq_{\text{cover}}$ does not satisfy *Improvement* by expanding the KB $\mathcal{K}$ in Example 4.3 with the formula $a \wedge b$.

As it turns out, contrary to the previous repair selection strategies, a key feature of the cover-based strategy is that the input KB is always covered by the selected repairs.

Proposition 4.18. *Let $\mathcal{K} \in \mathbb{K}$. If $R \cap R' = \emptyset$ for any $R, R' \in \mathcal{R}(\mathcal{K})$, then $\mathcal{F}(\mathcal{K}, \emptyset, \succcurlyeq_{\text{cover}}) = \mathcal{R}(\mathcal{K})$.*

The next example shows that even if the conflicts are binary, which makes $\succcurlyeq_{\text{comp}}$ strategy trivial by Proposition 4.12, the $\succcurlyeq_{\text{cover}}$ strategy is still able to give non-trivial results.

Example 4.19. Let $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ with $\mathcal{T} = \{A \sqsubseteq E, D \sqsubseteq \neg F, C \sqsubseteq F, C \sqsubseteq A, D \sqsubseteq \neg A, B \sqsubseteq \neg F\}$ and $\mathcal{A} = \{A(a), B(a), C(a), D(a)\}$. The conflicts of $\mathcal{K}$ are binary, i.e., $C(\mathcal{K}) = \{\{A(a), D(a)\}, \{B(a), C(a)\}, \{C(a), D(a)\}\}$. Then, $\mathcal{K}$ has three repairs, namely $R_1 = \{A(a), B(a)\}$, $R_2 = \{A(a), C(a)\}$, and $R_3 = \{B(a), D(a)\}$. Then, $\text{cover}(\mathcal{K}) = \{\{R_2, R_3\}\}$, thus $\mathcal{F}(\mathcal{K}, \emptyset, \succcurlyeq_{\text{cover}}) = \{R_2, R_3\} \subset \mathcal{R}(\mathcal{K})$.

4.2 *ch*-based Selection Strategies

Now, we focus on the second category of repair selection strategies that can take advantage of qualitative or quantitative information attached to formulas by the formula characterisation *ch*. For example, if the characterisation of formulas is numeric, i.e. $ch(\alpha) \in \mathbb{R}$, then we can apply any numeric aggregation function, such as summation, minimum, etc. In this work, besides this unary characterisation, we also consider binary relations, binary functions, and vector representations to characterise formulas in the input KB.

4.2.1 Improvement-based strategies. In the database and DL-Lite settings, the conflict is always binary, and the priority relation is suited for formulas forming a conflict. Here, we extend three existing formula characterisation-based methods —studied by the Database community [30] and applied to DL-Lite [7]— to an arbitrary binary relation $\succeq$ (without imposing any restriction on $\succeq$). This leads to the following improvement-based strategy defined through the repair selection strategies $\succcurlyeq_g$, $\succcurlyeq_p$ and $\succcurlyeq_c$ defined below:

Definition 4.20. Let $\mathcal{K} \in \mathbb{K}$, and $R, R' \in \mathcal{R}(\mathcal{K})$. We say that:

- R is a *global improvement* of R' w.r.t. $\succeq$, written $R \succcurlyeq_g R'$, if $\forall \beta \in R' \setminus R$, $\exists \alpha \in R \setminus R'$ such that $\alpha \succ \beta$.
- R is a *pareto improvement* of R' w.r.t. $\succeq$, written $R \succcurlyeq_p R'$, if $\exists \alpha \in R \setminus R'$ such that $\alpha \succ \beta, \forall \beta \in R' \setminus R$.
- R is a *completion improvement* of R' w.r.t. $\succeq$, written $R \succcurlyeq_c R'$, if $R \succcurlyeq_g R'$ holds w.r.t. a completion $\succ^c$ of $\succ$.

We call that $R \in \mathcal{F}(\mathcal{K}, \succeq, \succcurlyeq_x)$ for $x \in \{g, p, c\}$ a *globally-optimal, pareto-optimal, completion-optimal repair.*

The global, pareto and completion improvement relations $\succcurlyeq_x$ ($x \in \{g, p, c\}$) are defined from the relations $\succeq$ of formulas. This is in line with the intuition of our compatibility-based strategy, but takes the formula ordering into account through the notion of elementary improvement.

Now, let us introduce the priority relation, written $\succ_{|}$, by restricting the strict order $\succ$ merely to the conflicts in a given KB. More formally, given $\alpha, \beta \in \mathcal{K}$, $\alpha \succ_{|} \beta$ iff $\alpha \succ \beta$ and there exists at least one conflict $C \in C(\mathcal{K})$ s.t. $\{\alpha, \beta\} \subseteq C$. The following result holds:

Proposition 4.21. *Consider a KB $\mathcal{K} \in \mathbb{K}$ and $R, R' \in \mathcal{R}(\mathcal{K})$. If R is a global (resp. pareto, completion) improvement of R' w.r.t. $\succ_{|}$, then R is a global (resp. pareto, completion) improvement of R' w.r.t. $\succeq$.*

It can be checked by Definition 4.20 that a repair $R \in \mathcal{F}(\mathcal{K}, \succeq, \succcurlyeq_x)$ implies that $R \in \mathcal{F}(\mathcal{K}, \succ_{|}, \succcurlyeq_x)$ for $x \in \{g, p, c\}$, where the latter is introduced in [7] for DL-Lite.

Table 1: Principles vs. repair selection strategies: $\surd$ (resp. $\surd^*$) means that the strategy satisfies the property (resp. under certain condition), $\otimes$ stands for unsatisfaction, and NA means inapplicability.

Principle	**Category I**				**Category II**						
	$\succcurlyeq_{\text{card}}$	Compatibility		$\succcurlyeq_{\text{cover}}$	Improvement			Distance			
		$\succcurlyeq^{\#}_{\text{comp}}$	$\succcurlyeq^{\supseteq}_{\text{comp}}$		$g_{\succ}$	$p_{\succ}$	$c_{\succ}$	$\succcurlyeq^{d_v}_{\text{dist}}$	$\succcurlyeq^{d_{\text{Hsdf}}}_{\text{dist}}$	$\succcurlyeq^{d_\infty}_{\text{dist}}$	$\succcurlyeq^{d_1}_{\text{dist}}$
Abstraction	$\surd$	$\surd$		$\surd$	$\surd$			$\surd$	$\surd$		
Equivalence Invariance	$\otimes$	$\surd$		$\surd$	$\surd$			$\otimes$	$\surd^*$		
Coverness	$\otimes$	$\otimes$		$\surd$	$\otimes$			$\otimes$	$\otimes$		
Non-Emptiness	$\surd$	$\surd$		$\surd$	$\surd$			$\surd$	$\surd$		
Monotonicity	$\surd$	$\surd$		$\surd$	$\surd$			NA	NA		
Non-Discrimination	$\surd$	$\surd$		$\surd$	$\surd$			$\surd$	$\surd$		
Reducibility	$\otimes$	$\otimes$		$\otimes$	$\otimes$			$\otimes$	$\otimes$		
Consistency	$\surd$	$\surd$		$\surd$	$\surd$			$\surd$	$\surd$		
Improvement	$\surd$	$\surd$ (DL-Lite), $\otimes$ (PL)		$\otimes$	$\surd$			$\surd$	$\otimes$		
Stability	$\otimes$	$\surd$ (DL-Lite), $\otimes$ (PL)		$\otimes$	$\surd$			$\otimes$	$\otimes$		
Uniqueness	$\otimes$	$\otimes$		$\otimes$	$\otimes$			$\otimes$	$\otimes$		
Non-Triviality	$\surd$	$\otimes$ (DL-Lite), $\surd$ (PL)		$\surd$	$\surd$			$\surd$	$\surd$		

Proposition 4.22. *$\succcurlyeq_g$, $\succcurlyeq_p$ and $\succcurlyeq_c$ satisfy Abstraction, Equivalence Invariance, Non-Emptiness, Monotonicity, Non-Discrimination, Consistency, Improvement, Stability, and Non-Triviality principles, except Coverness, Reducibility and Uniqueness.*

Note that the *Improvement* and *Stability* properties hold due to the fact that if $R \succcurlyeq_x R'$ ($x \in \{g, p, c\}$), then $R \cup \{\alpha\} \succcurlyeq_x R'$, where $\alpha \in \mathsf{Form}(\mathcal{L})$. Satisfying these two principles shows the interestingness of the improvement-based strategies.

4.2.2 Distance-based strategy. Another way to compare repairs is by their distances to the original KB[3]. The basic intuition here is that the closer a repair $R \in \mathcal{R}(\mathcal{K})$ is to $\mathcal{K}$, the better R is. This leads to a new repair selection strategy given below:

Definition 4.23. Let $\mathcal{K} \in \mathbb{K}$ and $R, R' \in \mathcal{R}(\mathcal{K})$. Given a distance $d(R, \mathcal{K})$ between a repair R and $\mathcal{K}$, we say that R is preferred to R', noted $R \succcurlyeq^{d}_{\text{dist}} R'$, iff $d(R, \mathcal{K}) \leq d(R', \mathcal{K})$.

This strategy takes into account the closeness between repairs and the KB for the optimal repair selection. We distinguish two ways for defining such a distance. The first one concerns the case where the formulas are attached with weights. The second is dedicated to the case where a distance between formulas is provided, which can be then lifted to a distance among sets of formulas.

A widely used distance is the **hamming distance** [14] between two sets, here a KB $\mathcal{K}$ and a repair $R \in \mathcal{R}(\mathcal{K})$, which returns the number of formulas on which R and $\mathcal{K}$ differ. We denote by $d_h(R, \mathcal{K}) = |\mathcal{K} \setminus R|$ the hamming distance between R and $\mathcal{K}$. When the formulas are not equally reliable, they can be associated with different values using a function $v : \mathcal{K} \mapsto \mathbb{R}^+$. Based on this, one can define the weighted hamming distance induced by v as follows:

Definition 4.24. For $\mathcal{K} \in \mathbb{K}$, the **weighted hamming distance** between a repair $R \in \mathcal{R}(\mathcal{K})$ and $\mathcal{K}$ is: $d_v(R, \mathcal{K}) = \sum_{\alpha \in \mathcal{K} \setminus R} v(\alpha)$.

Observe that minimising the distance from a repair $R \in \mathcal{R}(\mathcal{K})$ to the KB $\mathcal{K}$ implies the maximization of the overall weight associated to R. We point out here that the weighted hamming distance-based strategy $\succcurlyeq^{d_v}_{\text{dist}}$ follows the *optimal subset repairs* principle that minimizes the number of facts deletion in Databases [23].

[3]The notion of distance was employed in repairing inconsistent databases (see [5]).

The next result shows that the repair selection strategy $\succcurlyeq^{d_v}_{\text{dist}}$ generalizes two existing repair selection methods: the cardinality-based strategy, the hamming distance-based strategy, the scoring function-based strategy $\geq_{\text{score}}$ [22], and the weight-based strategy $\geq_w$ [9], where the last two strategies are defined below:

- $R \geq_{\text{score}} R'$ if $\sum_{\alpha \in R} |\mathcal{R}(\alpha, \mathcal{K})| \geq \sum_{\alpha \in R'} |\mathcal{R}(\alpha, \mathcal{K})|$.
- $R \geq_w R'$ if $\sum_{\alpha \in R} w(\alpha) \geq \sum_{\alpha \in R'} w(\alpha)$.

Proposition 4.25. *Let $\mathcal{K} \in \mathbb{K}$. The following results hold:*

- *If $v(\alpha) = v(\beta)$ for any $\alpha, \beta \in \mathcal{K}$, $\succcurlyeq^{d_v}_{\text{dist}} = \succcurlyeq^{d_h}_{\text{dist}} = \succcurlyeq_{\text{card}}$.*
- *If $v(\alpha) = |\mathcal{R}(\alpha, \mathcal{K})|$, then $\succcurlyeq^{d_v}_{\text{dist}} = \succcurlyeq_{\text{score}}$.*
- *If $v(\alpha) = w(\alpha)$, then $\succcurlyeq^{d_v}_{\text{dist}} = \geq_w$.*

Proposition 4.26. *$\succcurlyeq^{d_v}_{\text{dist}}$ satisfies Abstraction, Non-Emptiness, Non-Discrimination, Consistency, Improvement, and Non-Triviality principles. However, $\succcurlyeq^{d_v}_{\text{dist}}$ does not satisfy Equivalence Invariance, Coverness, Reducibility, Stability, or Uniqueness. Note that the Monotonicity is not applicable.*

Now, we refer to the case where there exists a distance between every pair of formulas, as we will see later. Then, such a distance can be lifted to a distance between subsets using for instance the *Hausdorff distance* [15]. In detail, the Hausdorff distance is a classical mathematical object to measure how far two subsets of a metric space are from each other.

Definition 4.27 ([29]). Let $\mathcal{K} \in \mathbb{K}$. We assume that each pair of formulas $\alpha, \beta \in \mathcal{K}$ is associated with a distance $d(\alpha, \beta) \in \mathbb{R}^+$. The **Hausdorff distance** between two sets $S_1, S_2 \subseteq \mathcal{K}$ is defined as:

$$d_{\text{Hsdf}}(S_1, S_2) = \max\Big(\sup_{\alpha \in S_1} \inf_{\beta \in S_2} d(\alpha, \beta), \sup_{\alpha \in S_2} \inf_{\beta \in S_1} d(\alpha, \beta)\Big).$$

By Definition 4.27, it is clear to see that for a given repair $R \in \mathcal{R}(\mathcal{K})$, $d_{\text{Hsdf}}(R, \mathcal{K}) = \sup_{\alpha \in \mathcal{K}} \inf_{\beta \in R} d(\alpha, \beta)$.

Proposition 4.28. *$\succcurlyeq^{d_{Hsdf}}_{\text{dist}}$ satisfies Abstraction, Non-Emptiness, Non-Discrimination, Consistency, and Non-Triviality principles. However, $\succcurlyeq^{d_{Hsdf}}_{\text{dist}}$ does not satisfy Coverness, Reducibility, Improvement, Stability, or Uniqueness. Note that the Monotonicity is not applicable. If $d(\alpha, \beta) = d(\alpha', \beta')$ for $\alpha \equiv \alpha'$ and $\beta \equiv \beta'$ s.t. $\{\alpha, \beta\} \subseteq \mathcal{K}$, $\{\alpha', \beta'\} \subseteq \mathcal{K}'$, then $\succcurlyeq^{d_{Hsdf}}_{\text{dist}}$ satisfies Equivalence Invariance.*

In some circumstances, a repair $R \in \mathcal{R}(\mathcal{K})$ can be represented as a vector by considering, for instance, the inconsistency measures of the formulas in R, e.g., [22, 32]. Using such a representation, we have several ways to define the distance between a repair and a KB.

Definition 4.29. Let $\mathcal{K} \in \mathbb{K}$ and $R \in \mathcal{R}(\mathcal{K})$. Let $\overrightarrow{R} = \langle v_1, \ldots, v_n \rangle$ and $\overrightarrow{\mathcal{K}} = \langle u_1, \ldots, u_n \rangle$ be the vector representations of R and $\mathcal{K}$, respectively[4]. We consider two **vector-based distances** between $\mathcal{K}$ and R defined as:

$$d_\infty(R, \mathcal{K}) = \|\overrightarrow{\mathcal{K}} - \overrightarrow{R}\|_\infty = \max_{1 \le i \le n} |v_i - u_i| \qquad (\infty\text{-norm})$$

$$d_1(R, \mathcal{K}) = \|\overrightarrow{\mathcal{K}} - \overrightarrow{R}\|_1 = \sum_{1 \le i \le n} |v_i - u_i| \qquad (1\text{-norm})$$

Proposition 4.30. *$\succcurlyeq^{d_\infty}_{\text{dist}}$ and $\succcurlyeq^{d_1}_{\text{dist}}$ satisfy Abstraction, Consistency, Non-Emptiness, Non-Discrimination, and Non-Triviality principles. However, $\succcurlyeq^{d_\infty}_{\text{dist}}$ and $\succcurlyeq^{d_1}_{\text{dist}}$ do not satisfy Coverness, Reducibility, Improvement, Stability, or Uniqueness. Note that the Monotonicity is not applicable. If $\overrightarrow{R} = \overrightarrow{R'}$ for $R \cong R'$, then $\succcurlyeq^{d_\infty}_{\text{dist}}$ and $\succcurlyeq^{d_1}_{\text{dist}}$ satisfy Equivalence Invariance.*

Proposition 4.31 shows that the (weighted) hamming distance is a special case of the 1-norm vector-based distance.

Proposition 4.31. *Let $\mathcal{K} \in \mathbb{K}$ s.t. $\mathcal{K} = \{\alpha_1, \ldots, \alpha_n\}$. Then:*

- *$d_h(R, \mathcal{K}) = d_1(R, \mathcal{K})$ by taking $\overrightarrow{\mathcal{K}} = \langle 1, \ldots, 1 \rangle$, $\overrightarrow{R} = \langle v_1, \ldots, v_n \rangle$ with $v_i = 1$ if $\alpha_i \in R$ and $v_i = 0$, otherwise.*
- *$d_v(R, \mathcal{K}) = d_1(R, \mathcal{K})$ by taking $\overrightarrow{\mathcal{K}} = \langle v(\alpha_1), \ldots, v(\alpha_n) \rangle$, $\overrightarrow{R} = \langle v_1, \ldots, v_n \rangle$ with $v_i = v(\alpha_i)$ if $\alpha_i \in R$ and $v_i = 0$, otherwise.*

In [32], the authors defined a preference relation $L^{\succeq}_{\max}$ among repairs, which compares the Top-1 elements of two vectors —sorted in decreasing order— composed by the inconsistency measures of the concerned formulas. This can be viewed as a special case of d_∞.

Proposition 4.32. *Let $\mathcal{K} \in \mathbb{K}$ and $R, R' \in \mathcal{R}(\mathcal{K})$. Then, $L^{\succeq}_{\max}(R', R)$ iff $d_\infty(R, \mathcal{K}) \ge d_\infty(R', \mathcal{K})$ for $\overrightarrow{\mathcal{K}} = \overrightarrow{0}$.*

As discussed above, our second category of repair selection strategies is based on certain formulas characterisation. In particular, such category can be instantiated in different ways. A basic method is to use inconsistency measures [1, 4, 16, 18, 26, 27] to rank order the formulas according to their contribution to make the KB inconsistent.

For the distance-based strategies, we can achieve it by the conflict-based partition approach. Formally, we say the **conflict-based partition** of $\mathcal{K} \in \mathbb{K}$ is $\mathcal{P}(\mathcal{K}) = (\mathcal{P}_1, \ldots, \mathcal{P}_m)$ with $m = |C(\mathcal{K})| + 1$ such that $\mathcal{P}_i = \{\alpha \in \mathcal{K} \mid |C(\alpha, \mathcal{K})| = i - 1\}$. One can use such a conflict-based partition of $\mathcal{K}$ to have a formula weight, formula distance function, and a vector representation of its repairs as follows:

- The conflict-based **weight** of $\alpha \in \mathcal{K}$ is $v(\alpha) = 1/i$, if $\alpha \in \mathcal{P}_i$.
- For $\alpha \in \mathcal{P}_i, \beta \in \mathcal{P}_j$ $(1 \le i, j \le m)$, the **level-based distance** between α and β, written $d_L(\alpha, \beta)$, is defined as (i) $d_L(\alpha, \beta) = |i - j| + 1$ if $\alpha \ne \beta$, and (ii) $d_L(\alpha, \alpha) = 0$.
- For a repair $R \in \mathcal{R}(\mathcal{K})$, the **partition-based repair vector** of R is $\overrightarrow{R} = \langle v_1, \ldots, v_m \rangle$ with $v_i = |R \bigcap \mathcal{P}_i|$ for $1 \le i \le m$.

[4] We assume that $\overrightarrow{R}$ and $\overrightarrow{\mathcal{K}}$ have the same dimension by (possibly) adding 0 for the missing values in R.

Example 4.33 (Example 4.3 contd.). $\mathcal{K}$ has two conflicts: $C_1 = \{a, \neg a\}$ and $C_2 = \{a, \neg b, \neg a \vee b\}$. So, $\mathcal{P}_C(\mathcal{K}) = (\emptyset, \{\neg a, \neg b, \neg a \vee b\}, \{a\})$. We have $v(a) = 1/3$, $v(\neg a) = v(\neg b) = v(\neg a \vee b) = 1/2$. Recall the three repairs $\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_3$ of $\mathcal{K}$ in Example 4.3. Then, $d_v(\mathcal{R}_1, \mathcal{K}) = 1/3$, $d_v(\mathcal{R}_2, \mathcal{K}) = d_v(\mathcal{R}_3, \mathcal{K}) = 1/2 + 1/2 = 1$. Hence, $\mathcal{R}_1$ is optimal w.r.t. $\succcurlyeq^{d_v}_{\text{dist}}$, coinciding with the $\succcurlyeq_{\text{card}}$ strategy.
For $\beta_1, \beta_2 \in \{\neg a, \neg b, \neg a \vee b\}$ and $\beta_1 \ne \beta_2$, we have $d_L(\beta_1, a) = 2$, $d_L(\beta_1, \beta_2) = 1$. So, $d_{\text{Hsdf}}(\mathcal{R}_1, \mathcal{K}) = 2$, $d_{\text{Hsdf}}(\mathcal{R}_2, \mathcal{K}) = d_{\text{Hsdf}}(\mathcal{R}_3, \mathcal{K}) = 1$. Hence, $\mathcal{R}_2, \mathcal{R}_3$ are the preferred repairs w.r.t. the $\succcurlyeq^{d_{\text{Hdsf}}}_{\text{dist}}$ strategy, which coincides with $\succcurlyeq^{\#}_{\text{comp}}$ (Example 4.6), as well as with $\succcurlyeq_{\text{cover}}$. Moreover, we have $\overrightarrow{\mathcal{R}_1} = \langle 0, 3, 0 \rangle$, $\overrightarrow{\mathcal{R}_2} = \overrightarrow{\mathcal{R}_3} = \langle 0, 1, 1 \rangle$, and $\overrightarrow{\mathcal{K}} = \langle 0, 3, 1 \rangle$. By the definition of ∞-norm and 1-norm, we have that $\mathcal{R}_1$ is the preferred repair w.r.t. $\succcurlyeq^{d_\infty}_{\text{dist}}$ and $\succcurlyeq^{d_1}_{\text{dist}}$ strategies.

Example 4.33 shows that the strategy $\succcurlyeq^{d_{\text{Hsdf}}}_{\text{dist}}$ yields a different result from other strategies (e.g., $\succcurlyeq^{d_\infty}_{\text{dist}}$ and $\succcurlyeq^{d_1}_{\text{dist}}$) though they share the same principles as described in Table 1. We conclude that satisfying the same principles does not guarantee agreement on optimal repair selection, which confirms the same conclusion for ranking of arguments in abstract argumentation frameworks [2]. In addition, we observe that $\succcurlyeq^{d_v}_{\text{dist}}$ and $\succcurlyeq^{d_\infty}_{\text{dist}}$ provide the same result for the KB $\mathcal{K}$. However, they differ from each other in terms of *Equivalence Invariance* and *Improvement* principles (see Table 1). It is clear that these two strategies cannot return the same result on all KBs. This shows that our principles give an abstract and effective way to discriminate among these optimal repair selection strategies.

5 FEASIBILITY STUDY

In this paper, we presented several novel strategies for selecting preferred repairs. Due to space constraints, we focus on a preliminary empirical study comparing the compatibility-based and the cardinality-based strategies. Given an inconsistent KB, our algorithm takes a set of conflicts as input and outputs its preferred repairs. This is a reasonable setting because the number of conflicts is often of limited size [6, 8, 32]. One main difference in our setting compared to existing literature is that the conflicts are not only binary as in the classical DL-Lite [6, 8]. To better implement the compatibility-based strategy, we mention the next two remarks:

Remark 1. For $R_1, R_2 \in \mathcal{R}(\mathcal{K})$, if R_1 and R_2 are compatible, there exists a conflict $C \in C(\mathcal{K})$ and $a, b, c \in C$ such that $a \in R_1 \setminus R_2$, $b \in R_2 \setminus R_1$, and $c \in R_1 \cap R_2$. Due to this fact, binary conflicts can be avoided while checking the compatibility of two repairs.

Remark 2. Consider the graph representation $G_C(\mathcal{K})$ of $\mathcal{K}$ [20]: Each vertex is labeled by a conflict $C \in C(\mathcal{K})$; For each $C, C' \in C(\mathcal{K})$ such that $C \cap C' \ne \emptyset$, there exists an edge between C and C'. Let $G_i (1 \le i \le n)$ be the connected components of $G_C(\mathcal{K})$ and $OR(\mathcal{K})$ be the preferred repairs of $\mathcal{K}$. Then, $OR(\mathcal{K}) = OR(G_1) \otimes \cdots \otimes OR(G_i) \otimes \cdots \otimes OR(G_n)$ (where $\mathbb{S}_1 \otimes \mathbb{S}_2 := \{S_1 \cup S_2 \mid S_1 \in \mathbb{S}_1, S_2 \in \mathbb{S}_2\}$ for a set S and $\mathbb{S}_1, \mathbb{S}_2 \subseteq 2^S$). So, the computation of preferred repairs of a KB can be decomposed into the computation of the preferred repairs of the connected components of $G_C(\mathcal{K})$.

To test the feasibility of our framework, we construct a KB using the National Downloadable File given by the Centers for Medicare and Medicaid Services[5], as outlined in [8]. The file contains data

[5] https://data.cms.gov/provider-data/

on healthcare professionals and their respective affiliations. Unlike [8], which only constructs binary conflicts, we also incorporated a query stipulating that the same location cannot have more than a certain number of physicians. This addition enables us to generate a substantial number of non-binary conflicts, with sizes ranging from 11 to 51. In total, we obtained 35710 (9134 binary and 26576 non-binary) conflicts for our experiments, resulting in 1888 connected components. The maximal/minimal/median/average number of formulas (resp. conflicts) in these components are 2469/2/17/25 (resp. 2438/1/6/20). The repairs are computed using the PySAT tool [19] for enumerating minimal hitting sets. All experiments were conducted on a laptop equipped with a 2-Core Intel core i5 CPU running at 2.70 GHz and 8 GB of RAM. The timeout was set to one minute for each connected component.

We have the following observations: (1) Among the 1888 components, 1588 (84.1%) succeeded in computing the $\succcurlyeq^{\#}_{\text{comp}}$ preferred repairs and all (100%) in computing the $\succcurlyeq_{\text{card}}$ preferred repairs. Moreover, the sum of the obtained $\succcurlyeq^{\#}_{\text{comp}}$ (resp. $\succcurlyeq_{\text{card}}$) preferred repairs is 10850 (resp. 36489), highlighting that the $\succcurlyeq^{\#}_{\text{comp}}$ strategy cuts off more repairs than the $\succcurlyeq_{\text{card}}$ one. (2) The $\succcurlyeq_{\text{card}}$ preferred repairs of the whole KB have a fixed size, which contains 32516 facts, while the $\succcurlyeq^{\#}_{\text{comp}}$ preferred repairs vary from 22895 facts (min) to 32516 facts (max) with the average equal to 31396. Therefore, there are cases where the $\succcurlyeq^{\#}_{\text{comp}}$ are also $\succcurlyeq_{\text{card}}$ preferred. (3) The total facts of our dataset is 46603, among which 34493 (resp. 33947) appear in some $\succcurlyeq^{\#}_{\text{comp}}$ (resp. $\succcurlyeq_{\text{card}}$) preferred repairs. This means that the $\succcurlyeq^{\#}_{\text{comp}}$ recovers more facts than $\succcurlyeq_{\text{card}}$. We conclude that computing preferred repairs are feasible for real-world KBs.

6 RELATED WORK

While a number of methods for selecting preferred repairs exist in knowledge and databases, this topic is still underdeveloped.

In [30], the authors proposed to select preferred repairs, under the improvement-based strategies, in relational databases based on acyclic preference relations among facts appearing in the same conflict (i.e., the facts that violate the same integrity constraint). Our framework considers the general case without imposing restriction on the preference relation, i.e., the preference ordering can hold between any pair of formulas.

Moreover, Staworko et al. studied some basic properties of repair selection functions [30], namely *Non-Emptiness*, *Monotonicity*, *Non-Discrimination*, *Categoricity* (a stricter version of *Uniqueness*), and *Conservativeness* (i.e., a preferred repair must be a repair). Note that *Conservativeness* is not relevant in our case since it is trivial. We enrich this principle set by introducing *Abstraction, Equivalence Invariance, Coverness, Consistency, Reducibility, Improvement, Stability,* and *Non-Triviality* principles. We show that the new principles allow us to discriminate different repair selection strategies. For example, $\succcurlyeq^{d_v}_{\text{dist}}$ behaves differently from the other distance-based strategies in terms of *Equivalence Invariance* and *Improvement*.

Another study which tackled the notion of preference among repairs in prioritized DL KBs was carried out by [7, 9]. In that setting, the authors proposed an inconsistency-tolerant framework to deal with query answering based on maximum cardinality repairs. Moreover, Bienvenu et al. applied the improvement-based strategies to filter preferred repairs from prioritized DL KBs [30]. In this paper, we identified two categories of repair selection strategies for arbitrary KBs, and studied their properties w.r.t. a set of 12 principles for PL and DL Lite.

To deal with conflicts over prioritized DL-Lite KBs, the authors in [6] studied methods for selecting only one preferred repair (w.r.t. priorities, deductive closure, cardinality and consistency criteria). However, the selected repair is not necessary $\subseteq$-maximal. Such limitation is not encountered in our framework since all the outputs of the presented selection strategies are indeed repairs. Indeed, *Uniqueness* contradicts with *Coverness* and *Conservativeness*, and is rejected by all the repair selection strategies studied in this paper.

Recently, [11] presented an inconsistency-tolerant query answering framework based on the selection from existential rules KBs the repairs that satisfy an $\subseteq$-maximal set of user preference rules. The authors have shown that their framework satisfies *Non-Emptiness*, *Non-Discrimination*, and *Conservativeness*. They also point out that if no priority relation among formulas is expressed, all repairs are selected as preferred repairs.

Note that the aforementioned works assume that the priority between formulas exists, which is not always the case in practice. In the absence of such priority, we provide two solutions. One is via selection strategies that require formula information, and the other is to discriminate among repairs by relying solely on the inherent structure of a KB.

Meanwhile, some closely related works used inconsistency measures [18, 31] to select the “less inconsistent" repairs, i.e., repairs containing formulas with less inconsistency values, in PL KBs [22] and in the ontology-based data access setting [32]. This is done by various aggregation functions, e.g. sum, max, lexmax, etc. Then, a set of desirable properties were considered to characterise the scoring and lifting functions. However, no principle has been designed to study the behaviour of the proposed preferred repair selection strategies. We showed that their strategies are encompassed by our framework. In particular, we proposed a family of principles for optimal repair selection strategies, instead of mere lifting functions. Indeed, to the best of our knowledge, this is the first principle-based framework handling the problem of selecting preferred repairs.

7 SUMMARY AND OUTLOOK

This paper introduced a principle-based framework to retrieve best repairs from inconsistent KBs under a broad family of strategies. Table 1 summarises the repair selection strategies, classified into two categories, and their satisfaction of a rich set of rationality principles. It is shown that our principles give an effective way to discriminate among these optimal repair selection strategies.

We plan to investigate more principles to analyse repair selection strategies in a fine-grained way, e.g. identifying the equivalence and the incompatibility among strategies [2]. We will deploy the strategies studied in this paper in the context of the well-known inconsistency-tolerant semantics. Moreover, we will study explanations for positive and negative query answers for the semantics under different optimal repair selection strategies [24].

ACKNOWLEDGMENTS

This work was funded, in part, by the Agence Nationale de la Recherche ANR under grant EXPIDA: ANR-22-CE23-0017.

REFERENCES

[1] Leila Amgoud and Jonathan Ben-Naim. 2015. Argumentation-based Ranking Logics. In *AAMAS*. 1511–1519.

[2] Leila Amgoud and Vivien Beuselinck. 2021. Equivalence of Semantics in Argumentation. In *International Conference on Principles of Knowledge Representation and Reasoning, KR*. 32–41.

[3] Leila Amgoud and Victor David. 2018. Measuring Similarity between Logical Arguments. In *KR*. 98–107.

[4] Meriem Ammoura, Yakoub Salhi, Brahim Oukacha, and Badran Raddaoui. 2017. On an MCS-based inconsistency measure. *Int. J. Approx. Reason.* 80 (2017), 443–459.

[5] Ofer Arieli, Marc Denecker, and Maurice Bruynooghe. 2007. Distance semantics for database repair. *Ann. Math. Artif. Intell.* 50, 3-4 (2007), 389–415.

[6] Salem Benferhat, Zied Bouraoui, and Karim Tabia. 2015. How to Select One Preferred Assertional-Based Repair from Inconsistent and Prioritized DL-Lite Knowledge Bases?. In *IJCAI*. 1450–1456.

[7] Meghyn Bienvenu and Camille Bourgaux. 2020. Querying and Repairing Inconsistent Prioritized Knowledge Bases: Complexity Analysis and Links with Abstract Argumentation. In *KR*. 141–151.

[8] Meghyn Bienvenu and Camille Bourgaux. 2022. Querying Inconsistent Prioritized Data with ORBITS: Algorithms, Implementation, and Experiments. In *KR*.

[9] Meghyn Bienvenu, Camille Bourgaux, and François Goasdoué. 2014. Querying Inconsistent Description Logic Knowledge Bases under Preferred Repair Semantics. In *AAAI*. 996–1002.

[10] Gerhard Brewka, Miroslaw Truszczynski, and Stefan Woltran. 2010. Representing Preferences Among Sets. In *AAAI*.

[11] Marco Calautti, Sergio Greco, Cristian Molinaro, and Irina Trubitsyna. 2020. Preference-based Inconsistency-Tolerant Query Answering under Existential Rules. In *KR*. 203–212.

[12] Marco Calautti, Sergio Greco, Cristian Molinaro, and Irina Trubitsyna. 2022. Preference-based inconsistency-tolerant query answering under existential rules. *Artif. Intell.* 312 (2022), 103772.

[13] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, and R. Rosati. 2007. Tractable Reasoning and Efficient Query Answering in Description Logics: The *DL-Lite* Family. *J. Autom. Reasoning* 39, 3 (2007), 385–429.

[14] Mukesh Dalal. 1988. Investigations into a Theory of Knowledge Base Revision. In *AAAI*, Howard E. Shrobe, Tom M. Mitchell, and Reid G. Smith (Eds.). 475–479.

[15] Michel C. Delfour and Jean-Paul Zolésio. 2011. *Shapes and Geometries - Metrics, Analysis, Differential Calculus, and Optimization*. Advances in design and control, Vol. 22.

[16] Jesse Heyninck, Badran Raddaoui, and Christian Straßer. 2023. Ranking-based Argumentation Semantics Applied to Logical Argumentation. In *IJCAI*. 3268–3276.

[17] Anthony Hunter and Sébastien Konieczny. 2008. Measuring Inconsistency through Minimal Inconsistent Sets. In *KR*. 358–366.

[18] Anthony Hunter and Sébastien Konieczny. 2010. On the measure of conflicts: Shapley Inconsistency Values. *Artif. Intell.* 174, 14 (2010), 1007–1026.

[19] Alexey Ignatiev, Antonio Morgado, and Joao Marques-Silva. 2018. PySAT: A Python Toolkit for Prototyping with SAT Oracles. In *SAT*. 428–437. https://doi.org/10.1007/978-3-319-94144-8_26

[20] Saïd Jabbour, Yue Ma, and Badran Raddaoui. 2014. Inconsistency measurement thanks to MUS decomposition. In *AAMAS*. 877–884.

[21] Saïd Jabbour and Badran Raddaoui. 2013. Measuring Inconsistency through Minimal Proofs. In *ECSQARU*. 290–301.

[22] Sébastien Konieczny, Pierre Marquis, and Srdjan Vesic. 2019. Rational Inference Relations from Maximal Consistent Subsets Selection. In *IJCAI*. 1749–1755.

[23] Ester Livshits, Benny Kimelfeld, and Sudeepa Roy. 2018. Computing Optimal Repairs for Functional Dependencies. In *PODS*. 225–237.

[24] Thomas Lukasiewicz, Enrico Malizia, and Cristian Molinaro. 2022. Explanations for Negative Query Answers under Inconsistency-Tolerant Semantics. In *IJCAI*. 2705–2711.

[25] Bernhard Nebel. 1991. Belief Revision and Default Reasoning: Syntax-Based Approaches. In *KR*. 417–428.

[26] Badran Raddaoui. 2018. On the measure of conflicts: an argumentation-based framework. *J. Appl. Non Class. Logics* 28, 2-3 (2018), 240–259.

[27] Badran Raddaoui, Christian Straßer, and Saïd Jabbour. 2023. A Comparative Study of Ranking Formulas Based on Consistency. In *IJCAI*. 3330–3337.

[28] Jandson S. Ribeiro and Matthias Thimm. 2021. Consolidation via Tacit Culpability Measures: Between Explicit and Implicit Degrees of Culpability. In *KR*. 529–538.

[29] R. Tyrrell Rockafellar and Roger J.-B. Wets. 1998. *Variational Analysis*. Springer Verlag.

[30] Slawek Staworko, Jan Chomicki, and Jerzy Marcinkowski. 2012. Prioritized repairing and consistent query answering in relational databases. *Ann. Math. Artif. Intell.* 64, 2-3 (2012), 209–246.

[31] Matthias Thimm and Johannes Peter Wallner. 2019. On the complexity of inconsistency measurement. *Artif. Intell.* 275 (2019), 411–456.

[32] Bruno Yun, Srdjan Vesic, Madalina Croitoru, and Pierre Bisquert. 2018. Inconsistency Measures for Repair Semantics in OBDA. In *IJCAI*, Jérôme Lang (Ed.). 1977–1983.