

Inconsistency-Tolerant Query Answering: Rationality Properties and Computational Complexity Analysis

Jean François Baget¹, Salem Benferhat², Zied Bouraoui^{3(✉)},
Madalina Croitoru⁴, Marie-Laure Mugnier⁴, Odile Papini⁵, Swan Rocher⁴,
and Karim Tabia²

¹ Inria, Paris, France

baget@lirmm.fr

² Artois University, Arras, France

{benferhat,tabia}@cril.fr

³ Cardiff University, Cardiff, UK

BouraouiZ@cardiff.ac.uk

⁴ Montpellier University, Montpellier, France

{croitoru,mugnier,rocher}@lirmm.fr

⁵ Aix-Marseille University, Marseille, France

odile.papini@lisis.org

Abstract. Generalising the state of the art, an inconsistency-tolerant semantics can be seen as a couple composed of a modifier operator and an inference strategy. In this paper we deepen the analysis of such general setting and focus on two aspects. First, we investigate the rationality properties of such semantics for existential rule knowledge bases. Second, we unfold the broad landscape of complexity results of inconsistency-tolerant semantics under a specific (yet expressive) subclass of existential rules.

1 Introduction

Within the Ontology-Based Data Access [17, 18] setting, this paper addresses the problem of query answering when the assertional base (which stores data) is inconsistent with the ontology (which represents generic knowledge about a domain). Recently, a general framework for inconsistency-tolerant semantics was proposed in [2]. This framework considers two key notions: modifiers and inference strategies. Inconsistency-tolerant query answering is seen as made out of a *modifier*, which transforms the original ABox into a so-called MBox, which is a set of consistent ABoxes (w.r.t. the TBox), and an *inference strategy*, which evaluates queries against this MBox knowledge base. Interestingly enough, such setting unifies main existing work and captures various semantics in the literature (see e.g., [1, 6, 16]). The obtained semantics were compared with respect to the productivity of their inference.

This paper goes one step further in the characterization of these inconsistency-tolerant semantics by carrying out an analysis in terms of rationality properties and data complexity. The rationality properties are considered for existential rule knowledge bases [3, 9] (a prominent ontology language that generalizes lightweight description logics). On the one hand we study basic properties of semantics such as their behaviour with respect to the conjunction and consistency of inferred conclusions. On the other hand, starting from the obvious observation that inconsistency-tolerant semantics are inherently nonmonotonic, we investigate their behaviour with respect to properties introduced for non-monotonic inference [14] that we rephrase in our framework. Entailment with general existential rules being undecidable, complexity is studied for a specific (yet expressive) subclass of existential rules known as *Finite Unification Sets* (FUS) [3], which in particular generalizes the description logic $DL\text{-Lite}_{\mathcal{R}}$ dedicated to query answering [10] (see also the *OWL2-QL* profile).

Before presenting our contributions, we provide some preliminaries on the logical setting and briefly recall the unified framework for inconsistency-tolerant semantics.

2 Preliminaries

We consider first-order logical languages without function symbols, hence a *term* is a variable or a constant. An *atom* is of the form $p(t_1, \dots, t_k)$ where p is a predicate name of arity k , and the t_i are terms. A (factual) *assertion* is an atom without variables (also named a *ground* atom). A *Boolean conjunctive query*¹ (and simply *query* in the following) is an existentially closed conjunction of atoms, that we will consider as a set of atoms, leaving quantifiers implicit. Given a set of assertions \mathcal{A} and a query q , the answer to q over \mathcal{A} is yes iff $\mathcal{A} \models q$, where \models denotes the standard logical consequence. Given two sets of atoms S_1 and S_2 (with disjoint sets of variables), a *homomorphism* h from S_1 to S_2 is a substitution of the variables in S_1 by the terms in S_2 such that $h(S_1) \subseteq S_2$ (where $h(S_1)$ is obtained from S_1 by substituting each variable according to h). It is well-known that, given two existentially closed conjunctions of atoms f_1 and f_2 (for instance queries and conjunctions of factual assertions), $f_1 \models f_2$ iff there is a homomorphism from the set of atoms in f_2 to the set of atoms in f_1 .

A knowledge base can be seen as a database enhanced with an ontological component. Since inconsistency-tolerant query answering has been mostly studied in the context of description logics (DLs), and especially DL-Lite, we will use some DL vocabulary, like ABox for the data and TBox for the ontology. However, our framework is not restricted to DLs, hence we define TBoxes and ABoxes in terms of first-order logic (and more precisely in the existential rule framework). We assume the reader familiar with the basics of DLs and their logical translation.

¹ For readability, we restrict our focus to *Boolean* conjunctive queries, however the framework and the obtained results can be directly extended to general conjunctive queries.

An *ABox* is a set of factual assertions. As a special case we have DL assertions restricted to unary and binary predicates. A *positive axiom* is of the form $\forall \mathbf{x} \forall \mathbf{y} (B[\mathbf{x}, \mathbf{y}] \rightarrow \exists \mathbf{z} H[\mathbf{y}, \mathbf{z}])$ where B and H are conjunctions of atoms; in other words, it is a *positive existential rule*. As a special case, we have for instance concept and role inclusions in DL-Lite $_{\mathcal{R}}$, which are respectively of the form $B_1 \sqsubseteq B_2$ and $S_1 \sqsubseteq S_2$, where $B_i := A \mid \exists S$ and $S_i := P \mid P^-$ (with A an atomic concept, P an atomic role and P^- the inverse of an atomic role). A *negative axiom* is of the form $\forall \mathbf{x} (B[\mathbf{x}] \rightarrow \perp)$ where B is a conjunction of atoms; in other words, it is a *negative constraint*. As a special case, we have for instance disjointness axioms in DL-Lite $_{\mathcal{R}}$, which are inclusions of the form $B_1 \sqsubseteq \neg B_2$ and $S_1 \sqsubseteq \neg S_2$, or equivalently $B_1 \cap B_2 \sqsubseteq \perp$ and $S_1 \cap S_2 \sqsubseteq \perp$.

A *TBox* $\mathcal{T} = \mathcal{T}_p \cup \mathcal{T}_n$ is partitioned into a set \mathcal{T}_p of positive axioms and a set \mathcal{T}_n of negative axioms. Finally, a *knowledge base* (KB) is of the form $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ where \mathcal{A} is an ABox and \mathcal{T} is a TBox. Such a KB is logically interpreted as the conjunction of its elements. \mathcal{K} is said to be *consistent* if $\mathcal{T} \cup \mathcal{A}$ is satisfiable, otherwise it is said to be *inconsistent*. We also say that \mathcal{A} is consistent (or inconsistent) with \mathcal{T} , which reflects the assumption that the TBox is reliable while the ABox may not. The answer to a query q over a consistent KB \mathcal{K} is yes iff $\langle \mathcal{T}, \mathcal{A} \rangle \models q$. When \mathcal{K} is inconsistent, standard consequence is not appropriate since all queries would be positively answered.

The notion of a (virtual) repair is a key notion in inconsistency-tolerant query answering. A repair is a subset of the ABox consistent with the TBox and inclusion-maximal for this property: $\mathcal{R} \subseteq \mathcal{A}$ is a *repair* of \mathcal{A} w.r.t. \mathcal{T} if (i) $\langle \mathcal{T}, \mathcal{R} \rangle$ is consistent, and (ii) $\forall \mathcal{R}' \subseteq \mathcal{A}$, if $\mathcal{R} \subsetneq \mathcal{R}'$ (\mathcal{R} is strictly included in \mathcal{R}') then $\langle \mathcal{T}, \mathcal{R}' \rangle$ is inconsistent. We denote by $\mathcal{R}(\mathcal{A})$ the set of \mathcal{A} 's repairs (for easier reading, we often leave \mathcal{T} implicit in our notations). Note that $\mathcal{R}(\mathcal{A}) = \{\mathcal{A}\}$ iff \mathcal{A} is consistent. The most commonly considered semantics for inconsistency-tolerant query answering, inspired from previous works in databases, is the following: q is said to be a *consistent consequence* of \mathcal{K} if it is a standard consequence of each repair of \mathcal{A} [1]. Several variants of this semantics have been proposed, which differ in their behaviour (cautiousness w.r.t. inconsistencies) and their computational complexity, see in particular [1, 6, 16].

3 A Unified Framework for Inconsistency-Tolerant Query Answering

In this section we recall the framework introduced in [2] for the study of inconsistency-tolerant query answering semantics. In this framework, semantics are defined by two components: a modifier and an inference strategy, applied on MBox knowledge bases. An *MBox KB* is simply a KB with multiple ABoxes of the form $\mathcal{K}_{\mathcal{M}} = \langle \mathcal{T}, \mathcal{M} \rangle$ where $\mathcal{T} = \mathcal{T}_p \cup \mathcal{T}_n$ is a TBox and $\mathcal{M} = \{\mathcal{A}_1, \dots, \mathcal{A}_m\}$ is a set of ABoxes, called an MBox. A standard KB will be seen as an MBox with $m = 1$. An MBox KB $\mathcal{K}_{\mathcal{M}}$ is said to be *consistent*, or \mathcal{M} is said to be consistent (with \mathcal{T}), if each \mathcal{A}_i in \mathcal{M} is consistent (with \mathcal{T}). A modifier transforms a possibly inconsistent MBox KB into an MBox KB such that, when the latter is

consistent, it can be provided as input to the inference strategy that determines if the query is entailed.

A (*composite*) *modifier* is a finite combination of elementary modifiers. In [2] the three following kinds of elementary modifiers are introduced:

- **Expansion** modifiers, which expand an MBox by explicitly adding some inferred assertions to its ABoxes. A natural expansion modifier is the *ground positive closure* of an MBox, which computes the closure of each ABox with respect to the positive axioms of the TBox, keeping only ground atoms:

$$\circ_{cl}(\mathcal{M}) = \{Cl(\mathcal{A}_i) \mid \mathcal{A}_i \in \mathcal{M}\}, \text{ where } Cl(\mathcal{A}_i) = \{\text{ground atom } a \mid \langle \mathcal{T}_p, \mathcal{A}_i \rangle \models a\}.$$

- **Splitting** modifiers, which replace each \mathcal{A}_i of an MBox by one or several of its maximally consistent subsets (hence, they always produce consistent MBoxes). A natural splitting modifier splits each ABox into the set of its *repairs*:

$$\circ_{rep}(\mathcal{M}) = \bigcup_{\mathcal{A}_i \in \mathcal{M}} \{\mathcal{R}(\mathcal{A}_i)\}.$$

- **Selection** modifiers, which select some elements of an MBox. A natural selection modifier is the *cardinality-based selection* modifier, which selects the largest ABoxes of an MBox:

$$\circ_{card}(\mathcal{M}) = \{\mathcal{A}_i \in \mathcal{M} \mid \forall \mathcal{A}_j \in \mathcal{M}, |\mathcal{A}_j| \leq |\mathcal{A}_i|\}.$$

Note that the cardinality-based selection function fully makes sense when inconsistency is due to the presence of multiple sources. Other selection functions, such as the ones based on rational closure or System Z [11] may be used, especially when inconsistency reflects the presence of exceptions in axioms of the TBox.

Many composite modifiers can be potentially defined using the three above “natural” modifiers, however this number is considerably reduced if we focus on non-equivalent modifiers: indeed, any composite modifier that produces a consistent MBox from a standard ABox, and obtained by combining the elementary modifiers \circ_{rep} , \circ_{card} and \circ_{cl} , is equivalent to one of the eight modifiers listed in Table 1. To ease reading, these modifiers are also denoted by abbreviations reflecting the order in which the elementary modifiers are applied, and using the following letters: R for \circ_{rep} , C for \circ_{cl} and M for \circ_{card} . Different kinds of inclusion relations hold between modifiers (see [2] for details).

Example 1. Let $\mathcal{K}_{\mathcal{M}} = \langle \mathcal{T}, \mathcal{M} \rangle$ be an MBox KB where $\mathcal{T} = \{A(x) \wedge B(x) \rightarrow \perp, A(x) \wedge C(x) \rightarrow \perp, B(x) \wedge C(x) \rightarrow \perp, A(x) \rightarrow D(x), B(x) \rightarrow D(x), C(x) \rightarrow D(x), B(x) \rightarrow E(x), C(x) \rightarrow E(x)\}$ and $\mathcal{M} = \{\{A(a), B(a), C(a), A(b)\}\}$. With R, we get $\circ_1(\mathcal{M}) = \{\{A(a), A(b)\}, \{B(a), A(b)\}, \{C(a), A(b)\}\}$. With CR: $\circ_5(\mathcal{M}) = \{\{A(a), D(a), A(b), D(b)\}, \{B(a), D(a), E(a), A(b), D(b)\}, \{C(a), D(a), E(a), A(b), D(b)\}\}$. With MCR: $\circ_6(\mathcal{M}) = \{\{B(a), D(a), E(a), A(b), D(b)\}, \{C(a), D(a), E(a), A(b), D(b)\}\}$.

An inference strategy takes as input a consistent MBox KB $\mathcal{K}_{\mathcal{M}} = \langle \mathcal{T}, \mathcal{M} \rangle$ and a query q and determines if q is entailed from $\mathcal{K}_{\mathcal{M}}$. Four main inference strategies


Table 1. The eight composite modifiers for an MBox $\mathcal{K}_{\mathcal{M}} = \langle \mathcal{T}, \mathcal{M} = \{\mathcal{A}\} \rangle$

Modifier	Combination	MBox
R	$\circ_1 = \circ_{rep}(\cdot)$	$\mathcal{M}_1 = \circ_1(\mathcal{M})$
MR	$\circ_2 = \circ_{card}(\circ_{rep}(\cdot))$	$\mathcal{M}_2 = \circ_2(\mathcal{M})$
CMR	$\circ_3 = \circ_{cl}(\circ_{card}(\circ_{rep}(\cdot)))$	$\mathcal{M}_3 = \circ_3(\mathcal{M})$
MCMR	$\circ_4 = \circ_{card}(\circ_{cl}(\circ_{card}(\circ_{rep}(\cdot))))$	$\mathcal{M}_4 = \circ_4(\mathcal{M})$
CR	$\circ_5 = \circ_{cl}(\circ_{rep}(\cdot))$	$\mathcal{M}_5 = \circ_5(\mathcal{M})$
MCR	$\circ_6 = \circ_{card}(\circ_{cl}(\circ_{rep}(\cdot)))$	$\mathcal{M}_6 = \circ_6(\mathcal{M})$
RC	$\circ_7 = \circ_{rep}(\circ_{cl}(\cdot))$	$\mathcal{M}_7 = \circ_7(\mathcal{M})$
MRC	$\circ_8 = \circ_{card}(\circ_{rep}(\circ_{cl}(\cdot)))$	$\mathcal{M}_8 = \circ_8(\mathcal{M})$


are considered, namely universal (also known as skeptical), safe, majority-based and existential (also called brave). They are formally defined as follows:

- *universal* consequence: $\mathcal{K}_{\mathcal{M}} \models_{\forall} q$ if $\forall \mathcal{A}_i \in \mathcal{M}, \langle \mathcal{T}, \mathcal{A}_i \rangle \models q$.
- *safe* consequence: $\mathcal{K}_{\mathcal{M}} \models_{\cap} q$ if $\langle \mathcal{T}, \bigcap_{\mathcal{A}_i \in \mathcal{M}} \mathcal{A}_i \rangle \models q$.
- *majority-based* consequence: $\mathcal{K}_{\mathcal{M}} \models_{maj} q$ if $\frac{|\{\mathcal{A}_i : \mathcal{A}_i \in \mathcal{M}, \langle \mathcal{T}, \mathcal{A}_i \rangle \models q\}|}{|\mathcal{M}|} > 1/2$.
- *existential* consequence: $\mathcal{K}_{\mathcal{M}} \models_{\exists} q$ if $\exists \mathcal{A}_i \in \mathcal{M}, \langle \mathcal{T}, \mathcal{A}_i \rangle \models q$.


Given two inference strategies s_i and s_j , s_i is said to be *more cautious* than s_j , denoted $s_i \leq s_j$, if for any consistent MBox $\mathcal{K}_{\mathcal{M}}$ and any query q , if $\mathcal{K}_{\mathcal{M}} \models_{s_i} q$ then $\mathcal{K}_{\mathcal{M}} \models_{s_j} q$. The considered inference strategies are totally ordered by \leq as follows: $\cap \leq \forall \leq maj \leq \exists$.




(a) Relationships between \cap -based semantics



(b) Relationships between \forall -based semantics



(c) Relationships between *maj*-based semantics



(d) Relationships between \exists -based semantics

Fig. 1. Productivity of inconsistency-tolerant semantics where $X \rightarrow Y$ means that Y is strictly more productive than X .

An inconsistency-tolerant query answering semantics is then defined by a composite modifier and an inference strategy.

Definition 1. Let $\mathcal{K}=\langle\mathcal{T},\mathcal{A}\rangle$ be a standard KB, \circ_i be a composite modifier and s_j be an inference strategy. A query q is said to be an $\langle\circ_i, s_j\rangle$ -consequence of \mathcal{K} , denoted by $\mathcal{K} \models_{\langle\circ_i, s_j\rangle} q$, if it is entailed from the MBox KB $\langle\mathcal{T}, \circ_i(\{\mathcal{A}\})\rangle$ by the strategy s_j .

Note that the main semantics from the literature [1, 6, 16] are covered by this definition: AR, IAR and ICR semantics respectively correspond to $\langle\mathbf{R}, \forall\rangle$, $\langle\mathbf{R}, \cap\rangle$, and $\langle\mathbf{CR}, \cap\rangle$.²

Example 2. Consider the input KB $\mathcal{K}_{\mathcal{M}}=\langle\mathcal{T}, \mathcal{M}\rangle$ from Example 1. $\circ_1(\mathcal{M}) = \mathcal{M}_1 = \{\{A(a), A(b)\}, \{B(a), A(b)\}, \{C(a), A(b)\}\}$. Since $A(b) \in \bigcap_{\mathcal{A}_i \in \mathcal{M}_1} \mathcal{A}_i$ and $A(x) \rightarrow D(x)$, $\mathcal{K} \models_{\langle\circ_1, \cap\rangle} D(b)$ holds. Hence, we also have $\mathcal{K} \models_{\langle\circ_1, \forall\rangle} D(b)$. Furthermore $\mathcal{K} \models_{\langle\circ_1, \forall\rangle} D(a)$. By $\langle\circ_1, maj\rangle$, $E(a)$ is furthermore entailed. Indeed, $\langle\mathcal{T}, \{B(a), A(b)\}\rangle \models E(a)$ and $\langle\mathcal{T}, \{C(a), A(b)\}\rangle \models E(a)$ and $|\mathcal{M}_1|=3$. By $\langle\circ_1, \exists\rangle$, $A(a)$ is also entailed. Let $q = \exists x D(x) \wedge E(x)$. Then q is a consequence of $\langle\circ_1, maj\rangle$ and $\langle\circ_1, \exists\rangle$.

The obtained semantics have been compared from a productivity point of view. Formally, a semantics $\langle\circ_i, s_k\rangle$ is *less productive* than a semantics $\langle\circ_j, s_l\rangle$ if, for any KB $\mathcal{K}=\langle\mathcal{T}, \mathcal{A}\rangle$ and any query q , if $\mathcal{K} \models_{\langle\circ_i, s_k\rangle} q$ then $\mathcal{K} \models_{\langle\circ_j, s_l\rangle} q$. This productivity relation is a preorder, which can be established by considering on the one hand the inclusion relations between composite modifiers and on the other hand the cautiousness total order on inference, as detailed below. Figure 1 depicts the results about semantics defined with the same inference strategy (note that transitivity edges are not drawn and no other edges hold). Then Theorem 1 extends these results to semantics possibly based on different inference strategies. In particular, if $s_k < s_l$ then, for all modifiers \circ_i and \circ_j , $\langle\circ_i, s_k\rangle$ is strictly less productive than $\langle\circ_i, s_l\rangle$, and $\langle\circ_j, s_l\rangle$ is at least as productive as $\langle\circ_i, s_k\rangle$.

Theorem 1 (Productivity of semantics [2]). *The inclusion relation \sqsubseteq is the smallest relation that contains the inclusions $\langle\circ_i, s_k\rangle \sqsubseteq \langle\circ_j, s_k\rangle$ defined by the inclusions in Fig. 1a to d and satisfying the two following conditions: (1) for all s_j, s_p and \circ_i , if $s_j \leq s_p$ then $\langle\circ_i, s_j\rangle \sqsubseteq \langle\circ_i, s_p\rangle$; (2) it is transitive.*

It follows from Theorem 1 that 26 different semantics are obtained (out of the possible 32 inference relations used in Fig. 1). We point out that this result holds even when KBs are restricted to DL-Lite \mathcal{R} TBoxes. Finally, note that when the initial KB is consistent, all semantics correspond to standard entailment, i.e., given a consistent standard KB \mathcal{K} and a query q , $\mathcal{K} \models_{\langle\circ_i, s\rangle} q$ iff $\mathcal{K} \models q$, for all $1 \leq i \leq 8$ and $s \in \{\cap, \forall, \exists, maj\}$.

² Note however that CAR and ICAR [16] are close to $\langle\mathbf{RC}, \forall\rangle$ and $\langle\mathbf{RC}, \cap\rangle$ resp., but not equivalent. They could be covered by considering other elementary modifiers.

4 Rationality Properties of Inconsistency-Tolerant Semantics

This section is dedicated to the logical properties of inconsistency-tolerant semantics. We first analyze the behaviour of these semantics w.r.t the conjunction (or set union) and the consistency of inferred conclusions for a fixed KB. We then turn our attention to the fact that these semantics are inherently non-monotonic. Indeed, if some query q is entailed from a KB using a semantics $\langle \circ_i, s_j \rangle$, then q may be questionable in the light of new factual assertions. We will assume that these new factual assertions are sure (and will speak of conditional inference, opposed to unconditional inference when the KB is fixed). Hence, we also analyze inconsistency-tolerant semantics w.r.t rationality properties introduced for nonmonotonic inference that we recast in our framework.

4.1 Properties of Unconditional Inference

Let $\mathcal{K}_{\mathcal{M}} = \langle \mathcal{T}, \{\mathcal{A}\} \rangle$ be a possibly inconsistent KB and $\langle \circ_i, s \rangle$ denote any semantics with $\circ_i \in \{\mathbf{R}, \mathbf{MR}, \mathbf{CMR}, \mathbf{MCMR}, \mathbf{CR}, \mathbf{MCR}, \mathbf{RC}, \mathbf{MRC}\}$ and $s \in \{\forall, \cap, \exists, maj\}$. We define the following desirable properties:

QCE (Query Conjunction Elimination) For any KB $\mathcal{K}_{\mathcal{M}}$ and any queries q_1 and q_2 , if $\mathcal{K}_{\mathcal{M}} \models_{\langle \circ_i, s \rangle} q_1 \wedge q_2$ then $\mathcal{K}_{\mathcal{M}} \models_{\langle \circ_i, s \rangle} q_1$ and $\mathcal{K}_{\mathcal{M}} \models_{\langle \circ_i, s \rangle} q_2$.

QCI (Query Conjunction Introduction) For any KB $\mathcal{K}_{\mathcal{M}}$ and any queries q_1 and q_2 , if $\mathcal{K}_{\mathcal{M}} \models_{\langle \circ_i, s \rangle} q_1$ and $\mathcal{K}_{\mathcal{M}} \models_{\langle \circ_i, s \rangle} q_2$ then $\mathcal{K}_{\mathcal{M}} \models_{\langle \circ_i, s \rangle} q_1 \wedge q_2$.

Cons (Consistency) For any set of assertions \mathcal{A}' , if $\mathcal{K}_{\mathcal{M}} \models_{\langle \circ_i, s \rangle} \mathcal{A}'$ then $\langle \mathcal{T}, \mathcal{A}' \rangle$ is consistent.

ConsC (Consistency of Conjunction) For any set of assertions \mathcal{A} , if for all $f \in \mathcal{A}$, $\mathcal{K}_{\mathcal{M}} \models_{\langle \circ_i, s \rangle} f$ then $\langle \mathcal{T}, \mathcal{A} \rangle$ is consistent.

ConsS (Consistency of Support) For any set of assertions \mathcal{A}' , if $\mathcal{K}_{\mathcal{M}} \models_{\langle \circ_i, s \rangle} \mathcal{A}'$ then there is $R \in \mathcal{R}(\mathcal{A})$, such that $\langle \mathcal{T}, R \rangle \models \mathcal{A}'$.

Note that in the three last properties, the sets of assertions could be extended to queries with a more complex formulation. We first remind that, when $\mathcal{K}_{\mathcal{M}}$ is consistent, all semantics correspond to standard entailment, hence $\mathcal{K}_{\mathcal{M}} \models_{\langle \circ_i, s \rangle} q_1 \wedge q_2$ iff $\mathcal{K}_{\mathcal{M}} \models_{\langle \circ_i, s \rangle} q_1$ and $\mathcal{K}_{\mathcal{M}} \models_{\langle \circ_i, s \rangle} q_2$. When $\mathcal{K}_{\mathcal{M}}$ is inconsistent, one direction is still true for all semantics, namely Property **QCE**, which relies on the consistency of a repair. The converse direction, namely Property **QCI**, is obviously satisfied by universal and safe semantics but not by brave and majority-based semantics, even when q_1 and q_2 are ground atoms and the TBox contains only disjointness inclusions as shown by the next examples.

*Example 3 (Majority-based semantics does not satisfy QCI).*³ Let $\mathcal{T} = \{B \sqcap C \sqsubseteq \perp, A \sqcap D \sqsubseteq \perp, C \sqcap D \sqsubseteq \perp\}$ and $\mathcal{A} = \{A(a), B(a), C(a), D(a)\}$. The repairs are $\{A(a), B(a)\}$, $\{A(a), C(a)\}$ and $\{B(a), D(a)\}$. All modifiers give the same MBox since $\mathcal{T}_p = \emptyset$ and the repairs have the same size. $A(a)$ and $B(a)$ are each entailed by a majority of repairs but their conjunction is not.

³ Most examples in this section are provided in DL-Lite \mathcal{R} in order to show that some rationality properties do not hold even in this simple fragment of existential rules.

*Example 4 (Brave semantics does not satisfy properties **QCI** and **ConsC**).* Let $\mathcal{T}=\{A \sqcap B \sqsubseteq \perp\}$ and $\mathcal{A}=\{A(a), B(a)\}$. The repairs are $\{A(a)\}$ and $\{B(a)\}$. All modifiers lead to the same MBox since $\mathcal{T}_p=\emptyset$ and the repairs have the same size. $A(a)$ and $B(a)$ are both brave consequences but their conjunction is not. Besides **ConsC** is not satisfied since $\langle \mathcal{T}, \{A(a), B(a)\} \rangle$ is inconsistent.

Property **Cons** is true for any semantics (again by the consistency of a repair). Property **ConsC** holds for universal and safe semantics, and is false for any brave semantics, even for $|A_j|=|A_k|=1$ and DL-Lite TBoxes restricted to disjointness inclusions (see Example 3). Majority-based semantics are an interesting case, since the expressivity of the ontological language plays a role: Property **ConsC** is satisfied by all majority-based semantics when the language is restricted to DL-Lite $_{\mathcal{R}}$ and not satisfied as soon as we allow concept inclusions of the form $A \sqcap B \sqsubseteq C$ or ternary disjointness axioms of the form $A \sqcap B \sqcap C \sqsubseteq \perp$, even with ground queries (see Example 6). The fundamental reason why majority-based semantics satisfy Property **ConsC** over DL-Lite $_{\mathcal{R}}$ KBs is that, in these KBs, conflicts (i.e., minimal inconsistent subsets of the ABox) are necessarily of size two. When two ground atoms a_1 and a_2 are inferred with a majority-based strategy, at least one of element of the considered (consistent) MBox classically entails both a_1 and a_2 , hence $a_1 \wedge a_2$ is consistent; when conflicts are of size two, pairwise consistency entails global consistency. Note that this property still holds if we extend DL-Lite $_{\mathcal{R}}$ to n-ary predicates.

*Example 5 (Majority-based semantics does not satisfy Property **ConsC** for slight generalizations of DL-Lite $_{\mathcal{R}}$).* Let $\mathcal{T}=\{A \sqcap B \sqcap C \sqsubseteq \perp\}$ and $\mathcal{A}=\{A(a), B(a), C(a)\}$. The repairs are $\{A(a), B(a)\}$, $\{A(a), C(a)\}$ and $\{B(a), C(a)\}$. All modifiers give the same MBox since $\mathcal{T}_p=\emptyset$ and all the repairs have the same size. Each atom from \mathcal{A} is entailed (by 2/3 repairs), however \mathcal{A} itself is not.

Finally, Property **ConsS**, which expresses that every conclusion has a consistent support in the ABox, is satisfied by all semantics except those involving modifiers RC and MRC (as illustrated by the next example).

*Example 6 ((M)RC-based semantics do not satisfy Property **ConsS**).*⁴ Let $\mathcal{T}=\{A \sqcap B \sqsubseteq \perp, A \sqsubseteq C_1, B \sqsubseteq C_2\}$ and $\mathcal{A}=\{A(a), B(a)\}$. The (maximal) repairs of the ABox' closure are $\{A(a), C_1(a), C_2(a)\}$ and $\{B(a), C_1(a), C_2(a)\}$. The set of atoms $A_j = \{C_1(a), C_2(a)\}$ is entailed by all semantics based on RC and MRC, however no consistent subset of \mathcal{A} allows to entail A_j using \mathcal{T} .

Proposition 1 (Properties of unconditional inference). *The behaviour of semantics $\langle \circ_i, s \rangle$, with $\circ_i \in \{\mathbf{R}, \mathbf{MR}, \mathbf{CMR}, \mathbf{MCMR}, \mathbf{CR}, \mathbf{MCR}, \mathbf{RC}, \mathbf{MRC}\}$ and $s \in \{\cap, \forall, \text{maj}, \exists\}$, with respect to Properties **QCE**, **QCI**, **Cons**, **ConsC** and **ConsS**, is stated in Table 2.*

⁴ This example also shows that CAR and ICAR [16] do not satisfy **ConsS** (although they do when the conclusion is a single atom).

Table 2. Properties of unconditional inferences.

Properties	$\langle \circ_i, \cap \rangle$	$\langle \circ_i, \forall \rangle$	$\langle \circ_i, Maj \rangle$	$\langle \circ_i, \exists \rangle$
QCE	✓	✓	✓	✓
QCI	✓	✓	×	×
Cons	✓	✓	✓	✓
ConsC	✓	✓	×	×
ConsS $\circ_i \in \{RC, MRC\}$	×	×	×	×
otherwise	✓	✓	✓	✓

*: Except for languages where conflict sets involve at most two elements, like DL-Lite_R

4.2 Properties of Conditional Inferences

We now analyze more finely the inconsistency-tolerant semantics by considering their properties in terms of nonmonotonic inference. Within propositional logic setting, several approaches have been proposed for nonmonotonic inference (e.g. [5, 12, 14]). In such approaches nonmonotonicity is essentially caused by the fact that initial knowledge used for inference process is incomplete, and thus, later information may come to enrich them, which generally leads to revise some of the a priori considered hypotheses.

Let $\mathcal{K}_{\mathcal{M}} = \langle \mathcal{T}, \{\mathcal{A}\} \rangle$ be a possibly inconsistent KB and $\mathcal{A}_{\alpha}, \mathcal{A}_{\beta}$ be two sets of assertions such that $\langle \mathcal{T}, \mathcal{A}_{\alpha} \rangle$ and $\langle \mathcal{T}, \mathcal{A}_{\beta} \rangle$ are consistent. Assume that \mathcal{A}_{α} is the newly added knowledge. Since \mathcal{A}_{α} is considered as more reliable than the assertions in the KB, we have to keep \mathcal{A}_{α} in every selected repair of the KB. For the sake of simplicity, we define the notion of the set of repairs of $\mathcal{K}_{\mathcal{M}}$ in presence of a new consistent set of assertions \mathcal{A}_{α} with respect to a modifier \circ_i : $\mathcal{M}_i^{\alpha} = \{R : R \in \circ_i(\{\mathcal{A} \cup \mathcal{A}_{\alpha}\}) \text{ and } \mathcal{A}_{\alpha} \subseteq R\}$. Now, we say that \mathcal{A}_{β} is a *nonmonotonic consequence* of \mathcal{A}_{α} w.r.t. $\mathcal{K}_{\mathcal{M}}$, denoted by $\mathcal{A}_{\alpha} \sim_{\circ_i, s} \mathcal{A}_{\beta}$, if $\langle \mathcal{T}, \mathcal{M}_i^{\alpha} \rangle \models_s \mathcal{A}_{\beta}$.

In this study, we focus on the situation where the considered conclusions are sets of assertions, which can also be seen as conjunctions of ground queries. We first rephrase within our framework some KLM rationality properties [14]. Let $\mathcal{A}_{\alpha}, \mathcal{A}_{\beta}$ and \mathcal{A}_{γ} be consistent sets of assertions w.r.t \mathcal{T} and \sim be an inference relation, the *KLM* logical properties that we consider are the following⁵.

R (Reflexivity) $\mathcal{A}_{\alpha} \sim \mathcal{A}_{\alpha}$.

LLE (Left Logical Equivalence) If $\langle \mathcal{T}, \mathcal{A}_{\alpha} \rangle \equiv \langle \mathcal{T}, \mathcal{A}_{\beta} \rangle$ and $\mathcal{A}_{\alpha} \sim \mathcal{A}_{\gamma}$ then $\mathcal{A}_{\beta} \sim \mathcal{A}_{\gamma}$.

RW (Right Weakening) If $\langle \mathcal{T}, \mathcal{A}_{\alpha} \rangle \models \langle \mathcal{T}, \mathcal{A}_{\beta} \rangle$ and $\mathcal{A}_{\gamma} \sim \mathcal{A}_{\alpha}$ then $\mathcal{A}_{\gamma} \sim \mathcal{A}_{\beta}$.

Cut If $\mathcal{A}_{\alpha} \sim \mathcal{A}_{\beta}$ and $\mathcal{A}_{\alpha} \cup \mathcal{A}_{\beta} \sim \mathcal{A}_{\gamma}$ then $\mathcal{A}_{\alpha} \sim \mathcal{A}_{\gamma}$.

⁵ We have adopted here a formulation close to the one of *KLM* logical properties, even at the cost of simplicity. For instance $\langle \mathcal{T}, \mathcal{A}_{\alpha} \rangle \models \langle \mathcal{T}, \mathcal{A}_{\beta} \rangle$ could have been simplified in $\langle \mathcal{T}, \mathcal{A}_{\alpha} \rangle \models \mathcal{A}_{\beta}$. We remind that \models and \equiv denote standard logical entailment and equivalence.

CM (Cautious Monotony) If $\mathcal{A}_\alpha \vdash \mathcal{A}_\beta$ and $\mathcal{A}_\alpha \vdash \mathcal{A}_\gamma$ then $\mathcal{A}_\alpha \cup \mathcal{A}_\beta \vdash \mathcal{A}_\gamma$.
And If $\mathcal{A}_\alpha \vdash \mathcal{A}_\beta$ and $\mathcal{A}_\alpha \vdash \mathcal{A}_\gamma$ then $\mathcal{A}_\alpha \vdash \mathcal{A}_\beta \cup \mathcal{A}_\gamma$.

R means that the additional assertions have to be a consequence of the inference relation. **LLE** expresses the fact that two equivalent sets of assertions have the same consequences. **RW** says that consequences of the plausible assertions are plausible assertions too. **Cut** expresses the fact that if a plausible consequence is as secure as the assumptions it is based on, then it may be added into the assumptions. **CM** expresses that learning new assertions that could be plausibly inferred should not invalidate previous consequences. **And** expresses that the conjunction of two plausible consequences is a plausible consequence. The first five properties correspond to the system C [14] while the **And** property is derived from the previous ones. Clearly the **And** property is closely related to the **QCI** property given in Sect. 4.2. Indeed when $\mathcal{A}_\alpha = \emptyset$ (empty set, no additional information) and if q_1 and q_2 used in **CQI** are sets of assertions then **And** is equivalent to **CQI**. We now give the properties of the inference relations.

Proposition 2 (Properties of conditional inference). *The behaviour of inference relations $\vdash_{\circ_i, s}$, with $\circ_i \in \{R, MR, CMR, MCMR, CR, MCR, RC, MRC\}$ and $s \in \{\cap, \forall, maj, \exists\}$, with respect to Properties **R**, **LLE**, **RW**, **Cut**, **CM**, **And**, is given in Table 3.*

Proof: [Sketch of proof]. Properties **R**, **LLE** and **RW** follow from the definition of \mathcal{M}_i^α . For $s \in \{\forall, \cap, \exists\}$ and for $\circ_i \in \{R, MR\}$ the satisfaction of Properties **Cut** and **CM** stems from the fact that $\forall R' \in \mathcal{M}_i^{\alpha \cup \beta}$ we have $R' = R \cup \mathcal{A}_\beta$ with $R \in \mathcal{M}_i^\alpha$. Moreover, for $\circ_i \in \{CMR, CR, RC, MRC\}$ the satisfaction of Properties **Cut** and **CM** holds due the fact that $\forall R' \in \mathcal{M}_i^{\alpha \cup \beta}$ we have $R' = R \cup C(\mathcal{A}_\beta)$ with $R \in \mathcal{M}_i^\alpha$. The following counter-examples show the non-satisfaction cases. \square

Example 7 ($\not\vdash_{\circ_i, s}$ with $\circ_i \in \{MCMR, MCR\}$ and $s \in \{\forall, \exists, \cap\}$ does not satisfy **Cut**). For **MCMR**: Let $\mathcal{T} = \{A \sqsubseteq \neg B, A \sqsubseteq \neg G, F \sqsubseteq \neg B, B \sqsubseteq C, C \sqsubseteq D, A \sqsubseteq E\}$, and $\mathcal{A} = \{A(a), B(a), F(a), G(a)\}$, $\mathcal{A}_\alpha = \emptyset$,

Table 3. Properties of conditional inferences.

Properties	$\vdash_{\circ_i, \forall}$	$\vdash_{\circ_i, \cap}$	$\vdash_{\circ_i, \exists}$	$\vdash_{\circ_i, maj}$
R	✓	✓	✓	✓
LLE	✓	✓	✓	✓
RW	✓	✓	✓	✓
Cut $\circ_i \in \{MCMR, MCR\}$	×	×	×	×
otherwise	✓	✓	✓	×
CM $\circ_i \in \{MCMR, MCR\}$	×	×	×	×
otherwise	✓	✓	×	×
And	✓	✓	×	×

$\mathcal{A}_\beta = \{C(a), D(a)\}$, $\mathcal{A}_\gamma = \{A(a)\}$. We have $\mathcal{M}_4^\alpha = \{\{B(a), G(a), C(a), D(a)\}\}$ and $\mathcal{M}_4^{\alpha \cup \beta} = \{\{A(a), F(a), C(a), D(a), E(a)\}\}$. Thus $\langle \mathcal{T}, \mathcal{M}_4^\alpha \rangle \models_{\forall} \mathcal{A}_\beta$ and $\langle \mathcal{T}, \mathcal{M}_4^{\alpha \cup \beta} \rangle \models_{\forall} \mathcal{A}_\gamma$ but $\langle \mathcal{T}, \mathcal{M}_4^\alpha \rangle \not\models_{\forall} \mathcal{A}_\gamma$. **Cut** is not satisfied even for $s \in \{\exists, \cap\}$.
MCR: Let $\mathcal{T} = \{A \sqsubseteq \neg B, F \sqsubseteq \neg B, B \sqsubseteq C, C \sqsubseteq D\}$, $\mathcal{A} = \{A(a), B(a), F(a)\}$, $\mathcal{A}_\alpha = \emptyset$, $\mathcal{A}_\beta = \{C(a), D(a)\}$, $\mathcal{A}_\gamma = \{A(a)\}$. We have $\mathcal{M}_6^\alpha = \{\{B(a), C(a), D(a)\}\}$, $\mathcal{M}_6^{\alpha \cup \beta} = \{\{A(a), F(a), C(a), D(a)\}\}$. Thus $\langle \mathcal{T}, \mathcal{M}_6^\alpha \rangle \models_{\forall} \mathcal{A}_\beta$ and $\langle \mathcal{T}, \mathcal{M}_6^{\alpha \cup \beta} \rangle \models_{\forall} \mathcal{A}_\gamma$ but $\langle \mathcal{T}, \mathcal{M}_6^\alpha \rangle \not\models_{\forall} \mathcal{A}_\gamma$. **Cut** is not satisfied either for $s \in \{\exists, \cap\}$.

Example 8 ($\models \sim_{\circ_i, s}$ with $\circ_i \in \{\text{MCMR}, \text{MCR}\}$ and $s \in \{\forall, \cap\}$ does not satisfy **CM**). Let $\mathcal{T} = \{A \sqsubseteq \neg B, B \sqsubseteq C\}$, and $\mathcal{A} = \{A(a), B(a)\}$, $\mathcal{A}_\alpha = \emptyset$, $\mathcal{A}_\beta = \{C(a)\}$, $\mathcal{A}_\gamma = \{B(a)\}$. We have $\mathcal{M}_4^\alpha = \mathcal{M}_6^\alpha = \{\{B(a), C(a)\}\}$, $\mathcal{M}_4^{\alpha \cup \beta} = \mathcal{M}_6^{\alpha \cup \beta} = \{\{A(a), C(a)\}, \{B(a), C(a)\}\}$. Thus $\langle \mathcal{T}, \mathcal{M}_4^\alpha \rangle \models_{\forall} \mathcal{A}_\beta$ and $\langle \mathcal{T}, \mathcal{M}_4^\alpha \rangle \models_{\forall} \mathcal{A}_\gamma$ but $\langle \mathcal{T}, \mathcal{M}_4^{\alpha \cup \beta} \rangle \not\models_{\forall} \mathcal{A}_\gamma$. Moreover, $\langle \mathcal{T}, \mathcal{M}_6^\alpha \rangle \models_{\forall} \mathcal{A}_\beta$ and $\langle \mathcal{T}, \mathcal{M}_6^\alpha \rangle \models_{\forall} \mathcal{A}_\gamma$ but $\langle \mathcal{T}, \mathcal{M}_6^{\alpha \cup \beta} \rangle \not\models_{\forall} \mathcal{A}_\gamma$. **CM** is not satisfied even for $s = \cap$.

Example 9 ($\models \sim_{\circ_i, \text{maj}}$ with any \circ_i does not satisfy **Cut**). For $i = 1$ (R), let $\mathcal{T} = \{A \sqsubseteq \neg B, A \sqsubseteq \neg C, A \sqsubseteq \neg D, B \sqsubseteq \neg D, C \sqsubseteq \neg D, A \sqsubseteq E, B \sqsubseteq E, C \sqsubseteq E, D \sqsubseteq \neg E, A \sqsubseteq G, B \sqsubseteq G\}$, and $\mathcal{A} = \{A(a), B(a), C(a), D(a)\}$, $\mathcal{A}_\alpha = \{F(a)\}$, $\mathcal{A}_\beta = \{E(a)\}$, $\mathcal{A}_\gamma = \{G(a)\}$. We have $\mathcal{M}_1^\alpha = \{\{A(a), F(a)\}, \{B(a), F(a)\}, \{C(a), F(a)\}, \{D(a), F(a)\}\}$, thus $\langle \mathcal{T}, \mathcal{M}_1^\alpha \rangle \models_{\text{maj}} \mathcal{A}_\beta$. Moreover, $\mathcal{M}_1^{\alpha \cup \beta} = \{\{A(a), F(a), E(a)\}, \{B(a), F(a), E(a)\}, \{C(a), F(a), E(a)\}\}$ and $\langle \mathcal{T}, \mathcal{M}_1^{\alpha \cup \beta} \rangle \models_{\text{maj}} \mathcal{A}_\gamma$, however $\langle \mathcal{T}, \mathcal{M}_1^\alpha \rangle \not\models_{\text{maj}} \mathcal{A}_\gamma$. **Cut** is not satisfied for any other \circ_i .

Example 10 ($\models \sim_{\circ_i, \exists}$ with any \circ_i does not satisfy **CM**). For $i=1$ (R): Let $\mathcal{T} = \{A \sqsubseteq \neg C, A \sqsubseteq B, C \sqsubseteq B, A \sqsubseteq D, C \sqsubseteq E, D \sqsubseteq \neg C, E \sqsubseteq \neg A\}$, and $\mathcal{A} = \{A(a), C(a)\}$, $\mathcal{A}_\alpha = \{B(a)\}$, $\mathcal{A}_\beta = \{D(a)\}$, $\mathcal{A}_\gamma = \{E(a)\}$. We have $\mathcal{M}_1^\alpha = \{\{A(a), B(a)\}, \{C(a), B(a)\}\}$, thus $\langle \mathcal{T}, \mathcal{M}_1^\alpha \rangle \models_{\exists} \mathcal{A}_\beta$ and $\langle \mathcal{T}, \mathcal{M}_1^\alpha \rangle \models_{\exists} \mathcal{A}_\gamma$. Moreover $\mathcal{M}_i^{\alpha \cup \beta} = \{\{A(a), B(a), D(a)\}\}$ and $\langle \mathcal{T}, \mathcal{M}_1^{\alpha \cup \beta} \rangle \not\models_{\exists} \mathcal{A}_\gamma$. **CM** is not satisfied for any other \circ_i .

Example 11 ($\models \sim_{\circ_i, \text{maj}}$ with any \circ_i does not satisfy **CM**). For $i=1$ (R): Let $\mathcal{T} = \{A \sqsubseteq \neg B, A \sqsubseteq \neg C, B \sqsubseteq \neg C, A \sqsubseteq D, B \sqsubseteq D, C \sqsubseteq D, A \sqsubseteq E, B \sqsubseteq E, C \sqsubseteq F, B \sqsubseteq F, A \sqsubseteq \neg F\}$, and $\mathcal{A} = \{A(a), B(a), C(a)\}$, $\mathcal{A}_\alpha = \{D(a)\}$, $\mathcal{A}_\beta = \{E(a)\}$, $\mathcal{A}_\gamma = \{F(a)\}$. We have $\mathcal{M}_1^\alpha = \{\{A(a), D(a)\}, \{B(a), D(a)\}, R_3 = \{C(a), D(a)\}\}$, thus $\langle \mathcal{T}, \mathcal{M}_1^\alpha \rangle \models_{\text{maj}} \mathcal{A}_\beta$. Moreover $\langle \mathcal{T}, \mathcal{M}_1^\alpha \rangle \models_{\text{maj}} \mathcal{A}_\gamma$. We have $\mathcal{M}_1^{\alpha \cup \beta} = \{\{A(a), D(a), E(a)\}, \{B(a), D(a), E(a)\}\}$, thus $\langle \mathcal{T}, \mathcal{M}_1^{\alpha \cup \beta} \rangle \not\models_{\text{maj}} \mathcal{A}_\gamma$. **CM** is not satisfied for any other \circ_i .

Example 12 ($\models \sim_{\circ_i, s}$ with any \circ_i and $s \in \{\exists, \text{maj}\}$ does not satisfy **And**). For $i = 1$ and $s = \exists$ (R): Let \mathcal{T} and \mathcal{A} from Example 10. $\langle \mathcal{T}, \mathcal{M}_1^\alpha \rangle \models_{\exists} \mathcal{A}_\beta$ and $\langle \mathcal{T}, \mathcal{M}_1^\alpha \rangle \models_{\exists} \mathcal{A}_\gamma$ but $\langle \mathcal{T}, \mathcal{M}_1^\alpha \rangle \not\models_{\exists} \mathcal{A}_\beta \cup \mathcal{A}_\gamma$. **And** is not satisfied for any other \circ_i . For $i = 1$ and $s = \text{maj}$ (R): Let \mathcal{T} and \mathcal{A} from Example 11. $\langle \mathcal{T}, \mathcal{M}_1^\alpha \rangle \models_{\text{maj}} \mathcal{A}_\beta$ and $\langle \mathcal{T}, \mathcal{M}_1^\alpha \rangle \models_{\text{maj}} \mathcal{A}_\gamma$. but $\langle \mathcal{T}, \mathcal{M}_1^\alpha \rangle \not\models_{\text{maj}} \mathcal{A}_\beta \cup \mathcal{A}_\gamma$. **And** is not satisfied for any other \circ_i .

From Table 3, one can see that, for the composite modifiers $\circ_i \in \{\text{R,MR,CMR,CR,RC, MRC}\}$, the semantics based on universal and safe consequence satisfy all the properties of the system C. In **LLE** \mathcal{A}_γ can be replaced by a Conjunctive Query (CQ), in **RW** \mathcal{A}_α (resp. \mathcal{A}_β) can be replaced by CQ and **And** \mathcal{A}_γ (resp. \mathcal{A}_β) can be replaced by a CQ.

5 Complexity of Inconsistency-Tolerant Query Answering

In this section we study the data complexity⁶ of CQ entailment under the various semantics for classes of TBoxes $\mathcal{T} = \mathcal{T}_p \cup \mathcal{T}_n$ that fulfill the following property: \mathcal{T}_p is a *Finite Unification Set* (FUS) of existential rules [3], while \mathcal{T}_n remains any set of negative constraints. A set of rules \mathcal{T}_p fulfills the *FUS* property when, for any CQ q , there exists a *finite* set of CQs \mathcal{Q} (called the set of *rewritings* of q) such that for any ABox \mathcal{A} , $\langle \mathcal{T}_p, \mathcal{A} \rangle \models q$ iff $\exists q_i \in \mathcal{Q}$ such that $\mathcal{A} \models q_i$; in other words, q can be rewritten into a union of CQs \mathcal{Q} , which allows to forget the rules. Since query rewriting does not depend on any ABox, CQ entailment has the same data complexity as the classical database problem, which is in the low complexity class AC_0 . Note also that when \mathcal{T}_p satisfies the *FUS* property, the consistency of a standard KB can be checked by rewriting the query \perp with \mathcal{T} (or equivalently, rewriting each body of a negative constraint with \mathcal{T}_p) and checking if one of the obtained rewritings is entailed by \mathcal{A} . Such TBoxes encompass DL-Lite $_{\mathcal{R}}$ TBoxes as well as more expressive classes of existential rules, e.g., linear and sticky [8,9]. All the following membership results apply to *FUS* rules, while all hardness results hold as soon as DL-Lite $_{\mathcal{R}}$ TBoxes are considered.

We first briefly recall the definition of the complexity classes that we use. The class $\Delta_2^P = P^{NP}$ refers to problems solvable in polynomial time by a deterministic Turing Machine provided with an NP oracle, and its subclass $\Theta_2^P = \Delta_2^P[O(\log n)]$ is allowed to make only logarithmically many calls to an NP oracle. A Probabilistic Turing Machine (PTM) is a non-deterministic TM allowed to “toss coins” to make decisions: we will use the Probabilistic Polynomial-time (*PP*) class that contains the problems solvable in polynomial time with probability strictly greater than $\frac{1}{2}$ by a PTM [13].⁷ We also recall that Δ_2^P , Θ_2^P and *PP* are all closed under complement. CQ entailment with DL-Lite $_{\mathcal{R}}$ TBoxes is *coNP-complete* under $\langle \text{R}, \forall \rangle$ and $\langle \text{RC}, \forall \rangle$ semantics, and in AC_0 under $\langle \text{R}, \cap \rangle$ and $\langle \text{RC}, \cap \rangle$ semantics (semantics respectively known as *AR*, *CAR*, *IAR* and *ICAR* [16]). It is *coNP-complete* under $\langle \text{CR}, \cap \rangle$ semantics (known as *ICR* [6]), and Θ_2^P -*complete* under $\langle \text{MR}, \forall \rangle$ and $\langle \text{MR}, \cap \rangle$ semantics [7]. We first show that these complexity results also hold for FUS existential rules.

Proposition 3. *If CQ-entailment under $\langle \text{R/RC/MR}, \forall \rangle$ and $\langle \text{R/RC/CR}, \cap \rangle$ belongs to some complexity class \mathcal{C} for DL-Lite $_{\mathcal{R}}$ TBoxes, then CQ-entailment remains in the same complexity class \mathcal{C} for the more general FUS existential rules.*

⁶ This complexity measure is usually considered for query answering problems. Only the data (here the ABox) are considered in the problem input.

⁷ PP includes NP, co-NP and Θ_2^P .

Table 4. Complexity: tight complexity results are in black font (completely new results marked by a star, the other being generalizations of known results to *FUS*). Membership results are in gray font.

Modifier	\cap	\forall	<i>Maj</i>	\exists
R	AC_0	$coNP-c$	$PP-c$ *	AC_0 *
MR	Θ_2^P-c	Θ_2^P-c	$PP^{NP[O(\log n)]}$	Θ_2^P
CMR	Θ_2^P	Θ_2^P-c *	$PP^{NP[O(\log n)]}$	Θ_2^P
MCMR	Θ_2^P	Θ_2^P-c *	$PP^{NP[O(\log n)]}$	Θ_2^P
CR	$coNP-c$	$coNP-c$	$PP-c$ *	AC_0 *
MCR	Θ_2^P	Θ_2^P-c *	$PP^{NP[O(\log n)]}$	Θ_2^P
RC	AC_0	$coNP-c$	PP	P
MRC	Θ_2^P	Θ_2^P-c *	$PP^{NP[O(\log n)]}$	Θ_2^P

Proof: [Sketch] Let us first consider $\langle R/RC, \forall \rangle$. One can obviously guess a repair \mathcal{R} and check in polynomial time (actually in AC_0) if $\langle \mathcal{T}, \mathcal{R} \rangle \models \perp$ (by rewriting all negative constraints and looking for a homomorphism from one of those rewritings into \mathcal{R}), and if $\langle \mathcal{T}, \mathcal{R} \rangle \not\models q$ via rewriting methods as well. Concerning $\langle MR, \forall \rangle$, the membership holds for any FUS rules for similar reasons, and by observing that one can compute the maximum size of a repair through logarithmically many calls to an NP oracle. For $\langle R/MRC, \cap \rangle$ the technique from [16] still holds; whereas for $\langle CR, \cap \rangle$, we guess a set of repairs $\mathcal{R} = \{\mathcal{R}_1, \dots, \mathcal{R}_k\}$, with k polynomially bounded by the number of homomorphisms from rewritings of the query q to $Cl(\mathcal{A})$, such that: for any homomorphism h from a rewriting q' of q to $Cl(\mathcal{A})$ there is $\mathcal{R}_i \in \mathcal{R}$ with $h(q') \not\subseteq \mathcal{R}_i$. There is a polynomial number of rewritings (for data complexity), hence a polynomial number of homomorphisms from these rewritings to the $Cl(\mathcal{A})$. \square

The previous observations explain the complexity results written in black font without star in Table 4. We now provide some new complexity results for other universal-based and existential-based semantics.

Proposition 4. *CQ entailment under $\langle R, \exists \rangle$ (hence $\langle CR, \exists \rangle$) is in AC_0 .*

Proof: [Sketch] We first compute a set \mathcal{Q} that contains all the rewritings of q with the rules from \mathcal{T}_p , as well as all their specialisations according to all possible partitions on terms. We also rewrite \perp (i.e., all negative constraints) into the set \mathcal{N} . We remove from \mathcal{Q} all rewritings q' such that an element of \mathcal{N} maps to q' by homomorphism. Finally, we add to each remaining rewriting $q'' \in \mathcal{Q}$ all inequalities between its terms, which yields \mathcal{Q}' . \mathcal{Q}' can be seen as a union of CQs with inequality predicates, hence a first-order query. We have that $\mathcal{K} \models_{\langle R, \exists \rangle} q$ iff $\mathcal{A} \models \mathcal{Q}'$. Therefore q is first-order rewritable w.r.t. \mathcal{T} , under $\langle R, \exists \rangle$ semantics. \square

Proposition 5. *For $\circ_i \in \{\text{CMR}, \text{MCMR}, \text{MCR}, \text{MRC}\}$, CQ entailment under $\langle \circ_i, \forall \rangle$ and $\langle \circ_i, \exists \rangle$ semantics is in Θ_2^P .*

Proof: [Sketch] Notice that we can compute the maximum size of a repair and the maximum size of the ground positive closure of a maximum-sized repair through logarithmically many calls to an NP oracle. Then with one more call to this oracle, we can check whether there is a repair \mathcal{R} that satisfies the cardinality constraints and such that $\langle \mathcal{T}, \mathcal{R} \rangle \not\models q$ (resp. $\langle \mathcal{T}, \mathcal{R} \rangle \models q$). Therefore, $\langle \circ_i, \forall \rangle$ (resp. $\langle \circ_i, \exists \rangle$) is in Θ_2^P . \square

Proposition 6. *For $\circ_i \in \{\text{CMR}, \text{MCMR}, \text{MCR}, \text{MRC}\}$, CQ entailment under $\langle \circ_i, \forall \rangle$ semantics is Θ_2^P -hard.*

Proof: We adapt the reduction from the problem ParitySAT built in [7] (which is a reduction to $\langle \text{MR}, \forall \rangle$ with an instance query). We “tweak” the query and the *TBox* so that the positive part of the *TBox* is empty; this ensures that $\langle \circ_i, \forall \rangle = \langle \circ_j, \forall \rangle$ for any $\circ_i, \circ_j \in \{\text{MR}, \text{CMR}, \text{MCMR}, \text{MCR}, \text{MRC}\}$. \square

For majority-based semantics, we rely on probabilistic algorithms and provide two completeness results, as stated by the next proposition.

Proposition 7. *Conjunctive Query entailment under $\langle \text{R}, \text{Maj} \rangle$ and $\langle \text{CR}, \text{Maj} \rangle$ semantics is PP-complete.*

Proof: [Sketch] Membership: We use the following algorithm: first choose a subset S of atoms from \mathcal{A} randomly, then if S is not a repair of \mathcal{K} , output NO with probability $\frac{1}{2}$. Otherwise (S is a repair), if $\langle \mathcal{T}, S \rangle \models q$, output NO with probability $\frac{1}{2^{n+1}}$; else ($\langle \mathcal{T}, S \rangle \not\models q$), output NO with probability 1. This procedure obviously runs in polynomial time and the idea is that each repair has the same probability of being selected in the first step ($\frac{1}{2^n}$), and by answering NO a few times when $\langle \mathcal{T}, S \rangle \models q$ we ensure that the algorithm will give the right answer with probability strictly greater than $\frac{1}{2}$.

Hardness: We consider the following problem coMajSAT: given a Boolean SAT formula, is the number of unsatisfying affectations strictly greater than half of all possible affectations? We recall that *PP* is closed under complement. We notice that the reduction from SAT to $\langle S, \forall \rangle$ built in [15], ensures that each repair corresponds exactly to an affectation of the SAT formula, and the obtained query q is evaluated to true iff there is at least one invalid affectation. Hence, the majority of affectations are invalid iff q is entailed by the majority of the repairs. Hence, this transformation yields a reduction from coMajSAT to $\langle \text{R}, \text{Maj} \rangle$. Since $\langle \text{R}, \text{Maj} \rangle = \langle \text{CR}, \text{Maj} \rangle$, the result also holds for $\langle \text{CR}, \text{Maj} \rangle$. \square

To further clarify the complexity picture, we give some complexity class membership results for the remaining semantics (Table 4, in gray font). CQ entailment under $\langle \text{RC}, \exists \rangle$ semantics is clearly in *P* since we can first compute the ground positive closure of the ABox in polynomial time and $\langle \text{R}, \exists \rangle$ is in AC_0 . For $\langle \text{MRC}, \text{Maj} \rangle$ semantics, the membership proof from Proposition 7 holds as soon as we have observed that we could first compute the ground positive closure of the ABox. For the remaining majority-based semantics, we use an argument similar to the one in Proposition 5 to show membership to $PP^{NP[O(\log n)]}$: we only need logarithmically many calls to an NP oracle to get the maximum cardinality of a repair. Concerning the remaining intersection-based semantics $\langle \circ_i, \cap \rangle$,

we observe that by calling independently the corresponding universal problem (\circ_i, \forall) on each atom from the ABox, we can build the intersection of all repairs, hence the Θ_2^P membership. Finally, an interesting question is to what extent pre-processing the data, independently from any query, can reduce the complexity of query entailment. It seems reasonable to require that the result of this pre-processing step takes space at most linear in the size of the data. For instance, let us consider $\langle MR, \forall \rangle$: if we precompute the maximum cardinality of a repair (stored in $\log_2(|\mathcal{A}|)$ space), the complexity of CQ entailment drops from Θ_2^P -c to *coNP*-c, i.e., the complexity of $\langle R, \forall \rangle$.

6 Concluding Remarks

The framework for inconsistency-tolerant query answering recently proposed in [2] covers some well-known semantics and introduces new ones. These semantics were compared with respect to productivity. We broaden the analysis by considering two other points of view. First, we initiate a study of rationality properties of inconsistency-tolerant semantics. Second, we complement known complexity results, on the one hand by extending them to the more general case of FUS existential rules, and on the other hand by providing tight complexity results on some newly considered semantics (computation of repairs or closed repairs with majority-based or brave inference, as well several cardinality-based modifiers with universal inference).

The most efficiently computable semantics are $\langle R, \cap \rangle$ and $\langle R, \exists \rangle$ (equal to $\langle CR, \exists \rangle$). The $\langle R, \cap \rangle$ semantics is the least productive semantics in the framework. However, if one considers the closure of the repairs to increase the productivity of $\langle R, \cap \rangle$, i.e., $\langle CR, \cap \rangle$, one obtains a semantics that computationally costs as the “natural” semantics $\langle R, \forall \rangle$. At the opposite, $\langle R, \exists \rangle$ may be considered as too adventurous and does not behave well from a rationality point of view since it produces conclusions that may be inconsistent with the ontology. More generally, universal and safe semantics satisfy the rationality properties for most modifiers, which is not the case of majority-based and existential semantics. In addition, for all semantics, RC and MRC, which compute the closure of an inconsistent ABox, may lead to consider as plausible a conclusion with a contestable support, and since they do not seem to bring any advantage compared to other semantics, they should be discarded. Despite majority-based semantics do not fulfil some desirable logical properties, they remain interesting for several reasons: they are only slightly more complex to compute than universal semantics (w.r.t. the same modifier) while being more productive, without being as adventurous as existential semantics. Hence, they may be considered as a good tradeoff between both semantics when the universal semantics appear to be insufficiently productive. We also recall that majority-based semantics behave better from a logical viewpoint when they are restricted to DL-Lite \mathcal{R} (and more generally, when the ontological language ensures that the size of the conflicts is at most two). Regarding the use of cardinality, cardinality-based modifiers can be used to counteract troublesome assertions that conflict with many others, however they behave strangely when the cardinality criterion is applied to closed repairs.

In summary, no semantics appears to outperform all the others in all of the considered criteria. Selecting a semantics means selecting a suitable trade-off between productivity (or, inversely, cautiousness), satisfaction of rationality properties and computational complexity. We believe that this choice depends on the applicative context.

In a future work, new semantics could be considered within the unified framework, like *no-objection* semantics [4]. Besides, the study of rationality properties could be extended to other properties, and the exact complexity of several semantics remains an open issue.

Acknowledgments. This work was supported by the projects ASPIQ (ANR-12-BS02-0003), PAGOGA (ANR-12-JS02-007-01) and the ERC Starting Grant 637277.

References

1. Arenas, M., Bertossi, L.E., Chomicki, J.: Consistent query answers in inconsistent databases. In: Proceedings of the Eighteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, pp. 68–79 (1999)
2. Baget, J.-F., Benferhat, S., Bouraoui, Z., Croitoru, M., Mugnier, M.-L., Papini, O., Rocher, S., Tabia, K.: A general modifier-based framework for inconsistency-tolerant query answering. In: Proceedings of the Fifteenth International Conference on Principles of Knowledge Representation and Reasoning, KR (2016)
3. Baget, J.-F., Leclère, M., Mugnier, M.-L., Salvat, E.: On rules with existential variables: Walking the decidability line. *Artif. Intell.* **175**(9–10), 1620–1654 (2011)
4. Benferhat, S., Bouraoui, Z., Croitoru, M., Papini, O., Tabia, K.: Non-objection inference for inconsistency-tolerant query answering. In: Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (2016)
5. Benferhat, S., Cayrol, C., Dubois, D., Lang, J., Prade, H.: Inconsistency management and prioritized syntax-based entailment. In: Proceedings of the 13th International Joint Conference on Artificial Intelligence, pp. 640–647 (1993)
6. Bienvenu, M.: On the complexity of consistent query answering in the presence of simple ontologies. In: Proceedings of the Twenty-Sixth Conference on Artificial Intelligence (2012)
7. Bienvenu, M., Bourgaux, C., Goasdoué, F.: Querying inconsistent description logic knowledge bases under preferred repair semantics. In: Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, pp. 996–1002 (2014)
8. Cali, A., Gottlob, G., Pieris, A.: Towards more expressive ontology languages: the query answering problem. *Artif. Intell.* **193**, 87–128 (2012)
9. Cali, A., Gottlob, G., Lukasiewicz, T.: A general datalog-based framework for tractable query answering over ontologies. *J. Web Sem.* **14**, 57–83 (2012)
10. Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M., Rosati, R.: Tractable reasoning and efficient query answering in description logics: the DL-Lite family. *J. Autom. Reason.* **39**(3), 385–429 (2007)
11. Casini, G., Straccia, U.: Rational closure for defeasible description logics. In: Janhunen, T., Niemelä, I. (eds.) JELIA 2010. LNCS (LNAI), vol. 6341, pp. 77–90. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-15675-5_9](https://doi.org/10.1007/978-3-642-15675-5_9)
12. Gärdenfors, P., Makinson, D.: Nonmonotonic inference based on expectations. *Artif. Intell.* **65**(2), 197–245 (1994)

13. Gill, J.: Computational complexity of probabilistic turing machines. *SIAM J. Comput.* **6**(4), 675–695 (1977)
14. Kraus, S., Lehmann, D.J., Magidor, M.: Nonmonotonic reasoning, preferential models and cumulative logics. *Artif. Intell.* **44**(1–2), 167–207 (1990)
15. Lembo, D., Lenzerini, M., Rosati, R., Ruzzi, M., Savo, D.F.: Inconsistency-tolerant semantics for description logics. In: Hitzler, P., Lukasiewicz, T. (eds.) *RR 2010. LNCS*, vol. 6333, pp. 103–117. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-15918-3_9](https://doi.org/10.1007/978-3-642-15918-3_9)
16. Lembo, D., Lenzerini, M., Rosati, R., Ruzzi, M., Savo, D.F.: Inconsistency-tolerant query answering in ontology-based data access. *J. Web Sem.* **33**, 3–29 (2015)
17. Lenzerini, M.: Ontology-based data management. In: *Proceedings of the 6th Alberto Mendelzon International Workshop on Foundations of Data Management 2012*, pp. 12–15 (2012)
18. Poggi, A., Lembo, D., Calvanese, D., De Giacomo, G., Lenzerini, M., Rosati, R.: Linking data to ontologies. *J. Data Semant.* **10**, 133–173 (2008)