A General Modifier-based Framework for Inconsistency-Tolerant Query Answering

Jean François Baget
INRIA, France
baget@lirmm.fr

Salem Benferhat
Univ Artois, France
benferhat@cril.fr

Zied Bouraoui
Univ Aix-Marseille, France
zied.bouraoui@lsis.fr

Madalina Croitoru
Univ Montpellier, France
croitoru@lirmm.fr

Marie-Laure Mugnier
Univ Montpellier, France
mugnier@lirmm.fr

Odile Papini
Univ Aix-Marseille, France
odile.papini@amu-univ.fr

Swan Rocher
Univ Montpellier, France
rocher@lirmm.fr

Karim Tabia
Univ Artois, France
tabia@cril.fr

Abstract

We propose a general framework for inconsistency-tolerant query answering within existential rule setting. This framework unifies the main semantics proposed by the state of art and introduces new ones based on cardinality and majority principles. It relies on two key notions: modifiers and inference strategies. An inconsistency-tolerant semantics is seen as a composite modifier plus an inference strategy. We compare the obtained semantics from a productivity point of view.

Introduction

In this paper we place ourselves in the context of Ontology-Based Data Access and we address the problem of query answering when the assertional base (which stores data) is inconsistent with the ontology (which represents generic knowledge about a domain). Existing work in this area studied different inconsistency-tolerant inference relations, called *semantics*, which consist of getting rid of inconsistency by first computing a set of consistent subsets of the assertional base, called *repairs*, that restore consistency w.r.t the ontology, then using them to perform query answering. Most of these proposals, inspired by database approaches *e.g.* (Arenas, Bertossi, and Chomicki 1999) or propositional logic approaches *e.g.* (Benferhat, Dubois, and Prade 1997), were introduced for the lightweight description logic DL-Lite *e.g.* (Lembo et al. 2015). Other description logics *e.g.* (Rosati 2011) or existential rule *e.g.* (Lukasiewicz et al. 2015) have also been considered. In this paper, we use existential rules *e.g.*(Baget et al. 2011) as ontology language that generalizes lightweight description logics.

The main contribution of this paper consists in setting up a general framework that unifies previous proposals and extends the state of the art with new semantics. The idea behind our framework is to distinguish between the way data assertions are virtually distributed (notion of modifiers) and inference strategies. An inconsistency-tolerant semantics is then naturally defined by a modifier and an inference strategy. We propose a classification of the productivity of hereby obtained semantics by sound and complete conditions relying on modifier inclusion and inference strategy order. The objective of framework is to establish a methodology for inconsistency handling which, by distinguishing between modifiers and strategies, allows not only to cover existing semantics, but also to easily define new ones, and to study different kinds of their properties. Detailed proofs can be found in the associated technical report (Baget et al. 2016).

Preliminaries

We consider first-order logical languages without functional symbols, hence a *term* is a variable or a constant. In the following, by *query*, we mean a Boolean conjunctive query, i.e., an existentially quantified conjunction of atoms (note that more general kinds of queries could be considered). Given a set of facts $\mathcal{A}$ (atoms without variables) and a query q, the answer to q over $\mathcal{A}$ is yes iff $\mathcal{A}\models q$, where $\models$ denotes standard entailment.

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. We assume the reader familiar with the basics of DLs and their logical translation. An *ABox* is a set of 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$_R$, which are respectively of the form $B_1\sqsubseteq B_2$ and $S_1\sqsubseteq S_2$, where B_i:=$A|\exists S$ and S_i:=$P|P^-$ (with A an atomic concept, P an atomic role and P^- the inverse of an atomic role). A *negative axiom* is of the form $\forall\mathbf{x}(B[\mathbf{x}]\rightarrow\bot)$ 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$_R$, which are of the form $B_1\sqsubseteq\neg B_2$ and $S_1\sqsubseteq\neg S_2$. 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. 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. $\mathcal{K}$ 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 (in)consistent (with $\mathcal{T}$), which reflects the assumption that $\mathcal{T}$ is reliable. 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 entailment is not appropriate since all queries would be positively answered.

A key notion in inconsistency-tolerant query answering is the one of a repair of the ABox w.r.t. the TBox. A *repair* is a subset of the ABox consistent with the TBox and inclusion-maximal for this property. 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. Some of inconsistency-tolerant semantics use the notion of positive closure of an ABox. The *positive closure* of $\mathcal{A}$ (w.r.t. $\mathcal{T}$), denoted by *Cl*($\mathcal{A}$), is obtained by adding to $\mathcal{A}$ all assertions (built on the individuals occurring in $\mathcal{A}$) that can be inferred using the positive axioms of the TBox, namely: *Cl*($\mathcal{A}$)=$\{A$ atom $\mid \langle \mathcal{T}_p, \mathcal{A}\rangle \models A$ and *terms*$(A) \subseteq$ *terms*$(\mathcal{A})\}$. Note that the set of atomic consequences of a KB $\mathcal{K}$=$\langle \mathcal{T}, \mathcal{A}\rangle$ may be infinite whereas the positive closure of $\mathcal{A}$ is always finite since it does not contain new terms. Note also that $\mathcal{A}$ is consistent (with $\mathcal{T}$) iff *Cl*($\mathcal{A}$) is consistent (with $\mathcal{T}$). We can now recall the most well-known inconsistency-tolerant semantics (Arenas, Bertossi, and Chomicki 1999; Lembo et al. 2015; Bienvenu 2012). The most commonly considered semantics for inconsistency-tolerant query answering, inspired from previous work in databases, is the following: q is said to be a *consistent consequence* (or AR-consequence) of $\mathcal{K}$ if it is a standard consequence of each repair of $\mathcal{A}$. Variants of this semantics have been proposed. The CAR-entailment that consider a query as valid if it can be entailed using repairs computed from closed ABox. IAR-entailment (*resp.* ICAR-entailment) that considers the intersection of all repairs (*resp.* repairs computed from closed ABox), the ICR-entailment that considers the intersection of closed repairs.

A Unified Framework for Inconsistency-Tolerant Query Answering

In this section, we define a unified framework for inconsistency-tolerant query answering based on two main concepts: modifiers and inference strategies.

Let us first introduce the notion of MBox KBs. While a standard KB has a single ABox, it is convenient for subsequent definitions to introduce KBs with multiple ABoxes. Formally, an *MBox KB* is of the form $\mathcal{K}_{\mathcal{M}}$=$\langle \mathcal{T}, \mathcal{M}\rangle$ where $\mathcal{T}$ is a TBox and $\mathcal{M}$=$\{\mathcal{A}_1,\ldots,\mathcal{A}_n\}$ is a set of ABoxes called an MBox. We say that $\mathcal{K}_{\mathcal{M}}$ is *consistent*, or $\mathcal{M}$ is consistent (with $\mathcal{T}$) if each $\mathcal{A}_i$ in $\mathcal{M}$ is consistent (with $\mathcal{T}$). In the following, we start with an MBox KB which is a possibly inconsistent standard KB (i.e. with a single ABox in $\mathcal{M}$) and produce a consistent MBox KB, in which each element reflects a virtual reparation of the initial ABox.

Elementary and Composite Modifiers

We first introduce three classes of elementary modifiers: expansion, splitting and selection. For each class, we consider a "natural" instantiation, namely *positive closure*, splitting into *repairs* and selecting the largest elements (i.e., maximal w.r.t. *cardinality*). Elementary modifiers can be combined to define *composite* modifiers. Given the three natural instantiations of these modifiers, we show that their combination yields exactly eight different composite modifiers.

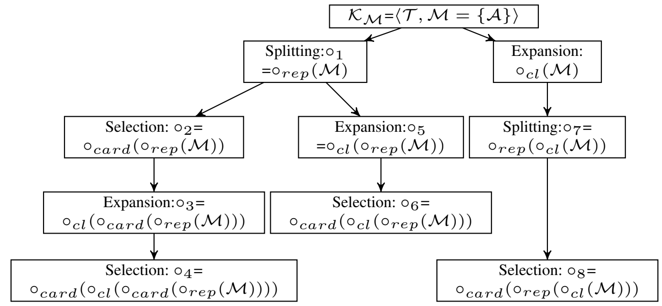

Figure 1: The eight possible combinations of modifiers from a single MBox KB $\mathcal{K}_{\mathcal{M}}$=$\langle \mathcal{T}, \mathcal{M} = \{\mathcal{A}\}\rangle$

Expansion modifiers. The expansion of an MBox consists in explicitly adding some inferred knowledge to its ABoxes. A natural expansion modifier consists in computing the *positive closure* of an MBox, which is defined as follows:

$$\circ_{cl}(\mathcal{M}) = \{Cl(\mathcal{A}_i) | \mathcal{A}_i \in \mathcal{M}\}.$$

Splitting modifiers. A splitting modifier always produces a consistent MBox and replaces each $\mathcal{A}_i$ of an MBox by one or several of its consistent subsets. A natural splitting modifier consists of splitting 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. A selection modifier selects some subsets of an MBox. As a natural selection modifier, we consider the *cardinality-based selection* modifier, which selects the largest elements of an MBox:

$$\circ_{card}(\mathcal{M}) = \{\mathcal{A}_i \in \mathcal{M} | \nexists \mathcal{A}_j \in \mathcal{M}\, s.t\, |\mathcal{A}_j| > |\mathcal{A}_i|\}.$$

We call a *composite modifier* any combination of these three elementary modifiers. We now study the question of how many different composite modifiers yielding consistent MBoxes exist and how do they compare with each other. We begin with some properties that considerably reduce the number of combinations to be considered. The three modifiers are idempotent and $\circ_{cl}$, $\circ_{rep}$ need to be applied once.

Lemma 1. *For any MBox $\mathcal{M}$, the following holds: (1) $\circ_{cl}(\circ_{cl}(\mathcal{M})) = \circ_{cl}(\mathcal{M})$, $\circ_{rep}(\circ_{rep}(\mathcal{M})) = \circ_{rep}(\mathcal{M})$ and $\circ_{card}(\circ_{card}(\mathcal{M})) = \circ_{card}(\mathcal{M})$. (2) Let $\circ_d$ be any composite modifier. Then $\circ_{cl}(\circ_d(\circ_{cl}(\mathcal{M}))) = \circ_d(\circ_{cl}(\mathcal{M}))$, and $\circ_{rep}(\circ_d(\circ_{rep}(\mathcal{M}))) = \circ_d(\circ_{rep}(\mathcal{M}))$.*

Figure 1 presents the eight different composite modifiers (thanks to Lemma 1) that can be applied to an MBox initially composed of a single (possibly inconsistent) ABox. At the beginning, one can perform either an expansion or a splitting operation (the selection has no effect). Expansion can only be followed by a splitting or a selection operation. After $\circ_{rep}(\circ_{cl}(\mathcal{M}))$ only a selection can be performed. Similarly, if one starts with a splitting operation followed by a selection operation, then only an expansion can be done. From $\circ_{cl}(\circ_{card}(\circ_{rep}(\mathcal{M})))$ only a selection can be performed.

To ease reading, we also denote the modifiers by short names reflecting the order in which the elementary modifiers are applied, using the following letters: R for $\circ_{rep}$, C for $\circ_{cl}$ and M for $\circ_{card}$ as shown in Table 1.

Modifier	Combination
R	$\circ_1 = \circ_{rep}(.)$
MR	$\circ_2 = \circ_{card}(\circ_{rep}(.))$
CMR	$\circ_3 = \circ_{cl}(\circ_{card}(\circ_{rep}(.)))$
MCMR	$\circ_4 = \circ_{card}(\circ_{cl}(\circ_{card}(\circ_{rep}(.))))$
CR	$\circ_5 = \circ_{cl}(\circ_{rep}(.))$
MCR	$\circ_6 = \circ_{card}(\circ_{cl}(\circ_{rep}(.)))$
RC	$\circ_7 = \circ_{rep}(\circ_{cl}(.))$
MRC	$\circ_8 = \circ_{card}(\circ_{rep}(\circ_{cl}(.)))$

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

Theorem 1. *Let $\mathcal{K}_{\mathcal{M}}$=$\langle\mathcal{T}, \mathcal{M} = \{\mathcal{A}\}\rangle$ be a possibly inconsistent KB. Then for any composite modifier $\circ_c$ that can be obtained by a finite combination of the elementary modifiers $\circ_{rep}$, $\circ_{card}$, $\circ_{cl}$, there exists a composite modifier $\circ_i$ in $\{\circ_1 \dots \circ_8\}$ (see Table 1) such that $\circ_c(\mathcal{M})$=$\circ_i(\mathcal{M})$.*

Example 1. *Let $\mathcal{K}_{\mathcal{M}}$=$\langle\mathcal{T}, \mathcal{M}\rangle$ be an MBox DL-Lite KB where $\mathcal{T}$=$\{A\sqsubseteq\neg B,\ A\sqsubseteq\neg C,\ B\sqsubseteq\neg C,\ A\sqsubseteq D,\ B\sqsubseteq D,\ C\sqsubseteq D,\ B\sqsubseteq E,\ C\sqsubseteq E\}$ and $\mathcal{M}$=$\{\{A(a), B(a), C(a), A(b)\}\}$. We have $\circ_1(\mathcal{M})$=$\{\{A(a), A(b)\}$,$\{B(a), A(b)\}$,$\{C(a), A(b)\}\}$, $\circ_5(\mathcal{M})$=$\{\{A(a), D(a), A(b), D(b)\}$,$\{B(a), D(a), E(a)$, $A(b)$,$D(b)\}$, $\{C(a), D(a), E(a), A(b), D(b)\}\}$, and $\circ_6(\mathcal{M})$=$\{\{B(a), D(a), E(a), A(b), D(b)\}$, $\{C(a), D(a)$, $E(a), A(b), D(b)\}\}$.*

The composite modifiers can be classified according to "inclusion" as depicted in Figure 2. We consider the relation, denoted $\subseteq_R$, defined as follows: given two modifiers X and Y, $X\subseteq_R Y$ if, for any MBox $\mathcal{M}$, for each $A\in X(\mathcal{M})$ there is $B\in Y(\mathcal{M})$ s.t $A\subseteq B$. We also consider two specializations of $\subseteq_R$: the "true" inclusion $\subseteq$ i.e. $X(\mathcal{M})\subseteq Y(\mathcal{M})$ and the "closure" inclusion, denoted $\subseteq_{cl}$: $X\subseteq_{cl}Y$ if $Y(\mathcal{M})$ is the closure of $X(\mathcal{M})$ (then each $A\in X(\mathcal{M})$ is included in its closure in $Y(\mathcal{M})$). In Figure 2, we label each edge by the most specific inclusion relation that holds from X to Y. Transitivity edges are not represented.

Inference Strategies for Querying an MBox

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 $\mathcal{K}_{\mathcal{M}}$. We consider four main inference strategies: universal, safe, majority-based and existential. We formally define these inference strategies as follows:

- Query q is a *universal* consequence of $\mathcal{K}_{\mathcal{M}}$, denoted by $\mathcal{K}_{\mathcal{M}} \models_{\forall} q$ iff $\forall\mathcal{A}_i \in \mathcal{M}$,$\langle\mathcal{T}, \mathcal{A}_i\rangle \models q$.
- Query q is a *safe* consequence of $\mathcal{K}_{\mathcal{M}}$, denoted by $\mathcal{K}_{\mathcal{M}} \models_{\cap} q$, iff $\left\langle\mathcal{T}, \bigcap_{\mathcal{A}_i\in\mathcal{M}} \mathcal{A}_i\right\rangle \models q$.
- Query q is a *majority-based* consequence of $\mathcal{K}_{\mathcal{M}}$, denoted by $\mathcal{K}_{\mathcal{M}} \models_{maj} q$, iff $\frac{|\mathcal{A}_i:\mathcal{A}_i\in\mathcal{M}, \langle\mathcal{T},\mathcal{A}_i\rangle\models q|}{|\mathcal{M}|} > 1/2$.
- Query q is an *existential* consequence of $\mathcal{K}_{\mathcal{M}}$, denoted by $\mathcal{K}_{\mathcal{M}} \models_{\exists} q$ iff $\exists\mathcal{A}_i \in \mathcal{M}$, $\langle\mathcal{T}, \mathcal{A}_i\rangle \models q$.

The *universal* inference strategy, also known as skeptical inference, is a standard way to derive conclusions from conflicting sources. It is used for instance in default reasoning where one only accepts conclusions derived from each extension of a default theory. The *safe* inference is a very sound and conservative inference relation since it only considers assertions shared by different ABoxes. The *existential* inference, also called brave inference, is a very adventurous inference relation and may derive conclusions that are together inconsistent with $\mathcal{T}$. It is often considered as undesirable when the KB represents available knowledge on some problem. It makes sense in some decision problems when one is only looking for a possible solution of a set of constraints or preferences. Finally, the *majority-based* inference considers as valid all conclusions entailed from $\mathcal{T}$ and the majority of ABoxes. It can be seen as a good compromise between universal/safe inference and existential inference.

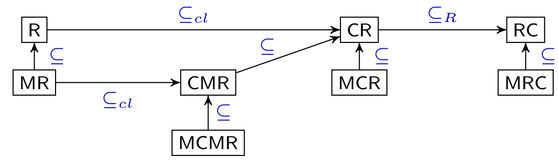

Figure 2: Inclusion relations between composite modifiers.

Given two inference strategies s_i and s_j, we say that s_i is *more cautious* than s_j, denoted $s_i \leq s_j$, when for any consistent $\mathcal{K}_{\mathcal{M}}$ 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 relations are totally ordered by $\leq$ as follows: $\cap \leq \forall \leq maj \leq \exists$

Example 2. *Let us consider the MBox $\mathcal{M}_1$=$\circ_1(\mathcal{M})$ given in Example 1. We have $\bigcap_{\mathcal{A}\in\mathcal{M}} \mathcal{A}_i$=$\{A(b)\}$, hence $\mathcal{K}_{\mathcal{M}_1} \models_{\cap} D(b)$. We also have $\mathcal{K}_{\mathcal{M}_1} \models_{\forall} D(a)$. The majority-based inference adds $E(a)$ as a valid conclusion. 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. Finally, we have $\mathcal{K}_{\mathcal{M}_1} \models_{\exists} A(a)$.*

Inconsistency-Tolerant Semantics = Composite Modifier + Inference Strategy

We can now define an inconsistency-tolerant semantics 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 $\langle\mathcal{T}, \circ_i(\{\mathcal{A}\})\rangle$ with the inference strategy s_j.*

Definition 1 covers the main semantics recalled in the preliminaries section: AR, IAR, CAR, ICAR and ICR respectively correspond to $\langle\circ_1,\forall\rangle$,$\langle\circ_1,\cap\rangle$,$\langle\circ_7,\forall\rangle$,$\langle\circ_7,\cap\rangle$ and $\langle\circ_5,\cap\rangle$.

Productivity Comparison of Inconsistency-Tolerant Semantics

We now compare the obtained semantics with respect to productivity, which we formalize as follows.

Definition 2. *Given two semantics $\langle\circ_i, s_k\rangle$ and $\langle\circ_j, s_l\rangle$, we say that $\langle\circ_j, s_l\rangle$ is* more productive *than $\langle\circ_i, s_k\rangle$, denoted $\langle\circ_i, s_k\rangle \sqsubseteq \langle\circ_j, s_l\rangle$, 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 $\mathcal{K} \models_{\langle\circ_j, s_l\rangle} q$.*

We first pairwise compare semantics defined with the same inference strategy. For each inference, we give necessary and sufficient conditions for the comparability of the associated semantics w.r.t. productivity. These conditions rely on the inclusion relations between modifiers (see Figure 2).

Proposition 1 (Productivity of $\cap$-semantics)**.** *See Figure 3. It holds that* $\langle \circ_i, \cap \rangle \sqsubseteq \langle \circ_j, \cap \rangle$ *iff* $\circ_j \subseteq \circ_i$ *or* $\circ_i \subseteq_R \circ_j$ *in a bijective way.*

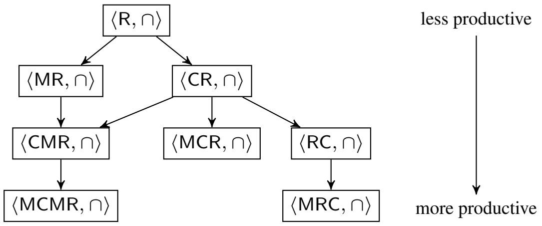

Figure 3: Relationships between $\cap$-based semantics

Proposition 2 (Productivity of $\forall$-semantics)**.** *See Figure 4. It holds that* $\langle \circ_i, \forall \rangle \sqsubseteq \langle \circ_j, \forall \rangle$ *iff* $\circ_j \subseteq \circ_i$ *or* $\circ_i \subseteq_R \circ_j$ *in a bijective way or* $\circ_j \subseteq_{cl} \circ_i$.

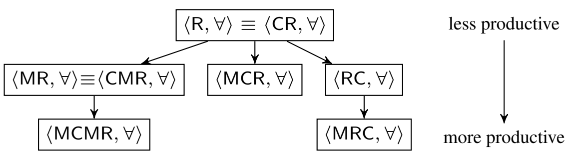

Figure 4: Relationships between $\forall$-based semantics

Proposition 3 (Productivity of maj-semantics)**.** *See Figure 5. It holds that* $\langle \circ_i, maj \rangle \sqsubseteq \langle \circ_j, maj \rangle$ *iff* $\circ_i \subseteq_R \circ_j$ *in a bijective way or* $\circ_j \subseteq_{cl} \circ_i$.

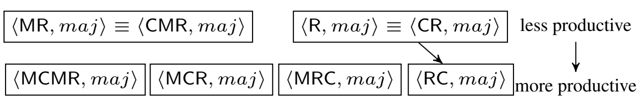

Figure 5: Relationships between maj-based semantics

Proposition 4 (Productivity of $\exists$-semantics)**.** *See Figure 6. It holds that* $\langle \circ_i, \exists \rangle \sqsubseteq \langle \circ_j, \exists \rangle$ *iff* $\circ_i \subseteq_R \circ_j$ *(in particular* $\circ_i \subseteq \circ_j$ *or* $\circ_i \subseteq_{cl} \circ_j$*) or* $\circ_j \subseteq_{cl} \circ_i$.

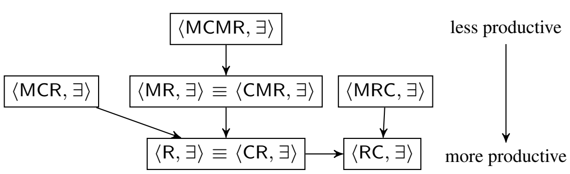

Figure 6: Relationships between $\exists$-based semantics

We now extend the previous results to any pair of semantics, possibly based on different inference strategies.

Theorem 2 (Productivity of semantics)**.** *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 Propositions 1-4 and satisfying the two following conditions: (1) for all* s_j*,* s_p *and* o_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.*

Theorem 2 is an important result. It states that the productivity relation can only be obtained from Propositions 1-4 and some composition of the relations. No more inclusion relations hold. In particular when $s_i > s_j$, it holds that $\forall k, \forall l$, $\langle \circ_k, s_i \rangle \not\sqsubseteq \langle \circ_l, s_j \rangle$, which means that there exist a query q and a KB $\mathcal{K}$ s.t q is an $\langle \circ_k, s_i \rangle$-consequence of $\mathcal{K}$ but not an $\langle \circ_l, s_j \rangle$-consequence of $\mathcal{K}$. Note that this holds already for DL-Lite$_R$ KBs. Lastly, note that when the initial KB is consistent, all semantics collapse with standard entailment.

Conclusion

This paper provides a general framework for inconsistency-tolerant query answering. On the one hand, our logical setting based on existential rules includes previously considered languages. On the other hand, viewing an inconsistency-tolerant semantics as a pair composed of a modifier and an inference strategy allows us to include the main known semantics and to consider new ones. We believe that the choice of semantics depends on the applicative context. In particular, cardinality-based selection allows us to counter troublesome assertions that conflict with many others. In some contexts, requiring to find an answer in all selected repairs can be too restrictive, hence the interest of majority-based semantics, which are more productive than universal semantics, without being as productive as the adventurous existential semantics. The productivity relations studied in this paper provided a criterion to compare different semantics. Rationality properties as well as complexity (which have been studied, but not presented in the paper due to the lack of space) provide other criteria for the choice of right inconsistency-tolerant semantics. As for future work, we plan consider other inference strategies such as the argued inference, parametrized inferences, etc. We also want to adapt the framework to belief change problems, like merging or revision.

Acknowledgment

This work has been supported by the French National Research Agency. ASPIQ project ANR-12-BS02-0003.

References

Arenas, M.; Bertossi, L. E.; and Chomicki, J. 1999. Consistent query answers in inconsistent databases. In *Proc. of SIGACT-SIGMOD-SIGART*, 68–79.

Baget, J.; Leclère, M.; Mugnier, M.; and Salvat, E. 2011. On rules with existential variables: Walking the decidability line. *Artif. Intell.* 175(9-10):1620–1654.

Baget, J. F.; Benferhat, S.; Bouraoui, Z.; Croitoru, M.; Mugnier, M.-L.; Papini, O.; Rocher, S.; and Tabia, K. 2016. A General Modifier-based Framework for Inconsistency-Tolerant Query Answering. *ArXiv e-prints*.

Benferhat, S.; Dubois, D.; and Prade, H. 1997. Some syntactic approaches to the handling of inconsistent knowledge bases: A comparative study part 1: The flat case. *Studia Logica* 58(1):17–45.

Bienvenu, M. 2012. On the complexity of consistent query answering in the presence of simple ontologies. In *Proc. of AAAI'12*.

Lembo, D.; Lenzerini, M.; Rosati, R.; Ruzzi, M.; and Savo, D. F. 2015. Inconsistency-tolerant query answering in ontology-based data access. *J. Web Sem.* 33:3–29.

Lukasiewicz, T.; Martinez, M. V.; Pieris, A.; and Simari, G. I. 2015. From classical to consistent query answering under existential rules. In *Proc. of AAAI'15*, 1546–1552.

Rosati, R. 2011. On the complexity of dealing with inconsistency in description logic ontologies. In *Proc. of IJCAI'11*, 1057–1062.