

Logic-Based Ranking of Assertions in Inconsistent ABoxes

Horacio Tellez Perez and Jef Wijsen

University of Mons, Belgium
{horacio.tellezperez, jef.wijsen}@umons.ac.be

Abstract. This paper concerns the problem of inconsistency in knowledge bases caused by ABox assertions. If such inconsistency occurs, users may be asked to rate the truthfulness of one or more ABox assertions. A potential difficulty is that users may not well understand how different ABox assertions interact (and possibly conflict) with one another through the TBox axioms. We therefore introduce a principled framework that uses TBox axioms for inferring positive and negative support for ABox assertions. This leads to a system of linear equations whose solution enhances the user’s truthfulness rating by means of logical information from the TBox.

1 Motivation

Inconsistency is an important and recurrent problem in today’s database and knowledge-base systems. Data errors are practically unavoidable in systems that integrate data coming from different sources. Two approaches to tackle this problem are data cleaning and consistent query answering (CQA) [Wij19]. The latter approach uses the notion of *repair*, which is a consistent knowledge base obtained by making some minimal amount of data corrections. A repair is conceptually not different from a cleaned knowledge base. However, whereas the process of data cleaning is supposed to end in a single cleaned knowledge base, the CQA approach allows for the possibility of multiple repairs (or possible worlds). In the Description Logics framework, different repairs correspond to different ABoxes, each one consistent with respect to a given fixed TBox. When we move from a single-world perspective to a possible-worlds perspective, different semantics for logical reasoning become possible [BR13,LLR⁺10,LB16,BBG14,GM12,CLP18]. Eventually, all these semantics address the following central problem: Given a logical sentence φ , assign some degree of truthfulness to φ . Is φ true in some, most, or all repairs? What is the probability of φ being true? Is there strong support for—or resistance against—the sentence φ ? Throughout this paper, we use the term “truthfulness” in a loose and informal way; we could have used other terms instead: credibility, veracity, accuracy...

In this paper, we present a new principled framework for evaluating the relative truthfulness of ABox assertions (i.e., the sentence φ is an ABox assertion in our framework). By relative, we mean that we are merely interested in ranking ABox assertions: Given two ABox assertions α and β , is α more, less, or equally truthful than β ? Our framework is different from CQA in that it does not use the notion of repair. The framework assumes that there is some initial weight function over the ABox assertions, called *credibility function*, which models the opinion of domain experts concerning the truthfulness of assertions. In their assessments, however, experts may have difficulties to fully understand and take into account the logic that is expressed by, possibly extensive, TBoxes. To overcome this difficulty, we propose an automated mathematical method for adjusting the experts' initial credibility function by taking into account the ontological logic of the TBox: strong logical support for an assertion α should increase its credibility, while strong logical support for $\neg\alpha$ should decrease α 's credibility. Eventually, our method results in a ranking of ABox assertions in terms of truthfulness, taking into account both the experts' credibility function and the TBox logic. In our framework, we assume that the TBox has been validated by a multitude of experts and is therefore error-free.

This paper is organized as follows. Section 2 presents a guiding example. Section 3 discusses related work. Section 4 introduces our general theoretical framework, and Section 5 presents a particular instantiation of this framework. Section 6 studies in more detail some theoretical questions and computational tasks raised by this instantiation. Finally, Section 7 concludes the paper.

2 Motivating Example

Consider the following knowledge base $\mathcal{K}_0 = \langle \mathcal{T}_0, \mathcal{A}_0 \rangle$.

$$\mathcal{T}_0 = \left\{ \begin{array}{l} \text{Professor} \sqsubseteq \text{Person}, \\ \text{Student} \sqsubseteq \text{Person}, \\ \text{Person} \sqsubseteq \neg\text{Course}, \\ \text{Student} \sqsubseteq \neg\text{Professor}, \\ \exists \text{teaches} \sqsubseteq \text{Professor}, \\ \exists \text{attends} \sqsubseteq \text{Student}, \\ \exists \text{teaches}^- \sqsubseteq \text{Course}, \\ \exists \text{attends}^- \sqsubseteq \text{Course} \end{array} \right\} \quad \mathcal{A}_0 = \left\{ \begin{array}{l} \text{John} : \text{Professor} \\ \text{Ava} : \text{Student} \\ \text{DB2} : \text{Course} \\ \text{KR} : \text{Course} \\ (\text{John}, \text{DB2}) : \text{teaches} \\ (\text{John}, \text{KR}) : \text{attends} \\ (\text{Ava}, \text{IA}) : \text{attends} \\ (\text{Bob}, \text{KR}) : \text{attends} \end{array} \right\}$$

This knowledge base is inconsistent, because from $(\text{John}, \text{KR}) : \text{attends}$, we can infer $\text{John} : \text{Student}$ by means of the axiom $\exists \text{attends} \sqsubseteq \text{Student}$. Then, by means of the axiom $\text{Student} \sqsubseteq \neg\text{Professor}$, we can infer $\text{John} : \neg\text{Professor}$, which obviously contradicts the first assertion in \mathcal{A}_0 . In conclusion, $(\text{John}, \text{KR}) : \text{attends}$ and $\text{John} : \text{Professor}$ contradict one another. The vertices in the directed graph of Fig. 1 are the assertions of \mathcal{A}_0 . A red-colored edge from α to β means that α *refutes* β . On the other hand, a green-colored edge from α to β means that α *supports* β . For example, from $(\text{John}, \text{DB2}) : \text{teaches}$ we can infer $\text{John} : \text{Professor}$

by means of the axiom $\exists teaches \sqsubseteq \text{Professor}$. Therefore, $(\text{John}, \text{DB2}) : teaches$ supports $\text{John} : \text{Professor}$. In this example, we assumed that domain experts esteemed that assertions in the ABox were equally credible, represented by a value of 1. Then, our proposed method updates values by balancing the value of each assertion α with respect to the values of α 's refuters and supporters. Refuters and supporters of α , respectively, lower and increase α 's value by an amount that is proportional to their own value. In Fig. 1, we see that $\text{John} : \text{Professor}$ is attributed a higher value than $(\text{John}, \text{KR}) : attends$ because it has more supporters and less refuters.

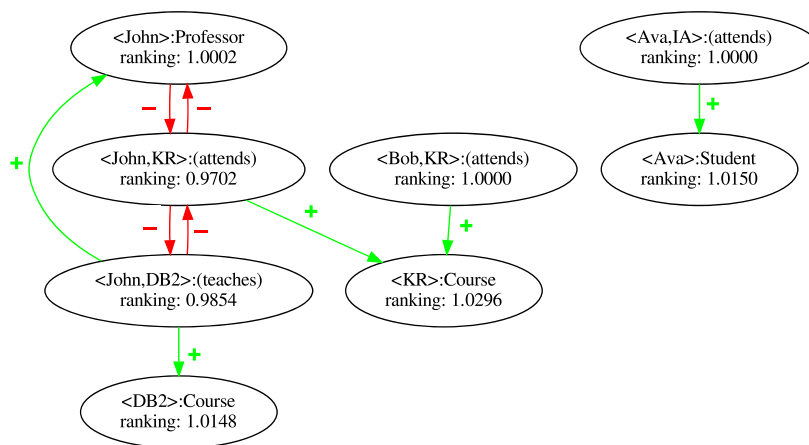


Fig. 1. ABox assessment obtained by a practical application of our framework.

Figure 2 shows the effect of adding $\text{Ava} : \text{Professor}$ to \mathcal{A}_0 . One can observe that the values of the assertions $(\text{Ava}, \text{IA}) : attends$ and $\text{Ava} : \text{Student}$ have gone down because of conflicts with the new assertion.

Since supporters and refuters are single assertions in this simple example, they can be represented in a directed graph. In our general framework, however, supporters and refuters will be *sets* of assertions. For example, if $A \sqcap B \sqsubseteq \neg C$ is a TBox assertion, then the set $\{(i, A), (i, B)\}$ is a refuter of (i, C) , but neither (i, A) nor (i, B) is a refuter on its own.

3 Related Work

In Description Logics, different repairs correspond to different ABoxes, each one consistent with respect to a given fixed TBox. When we move from a single ABox to a set of possible ABoxes, different semantics for logical reasoning

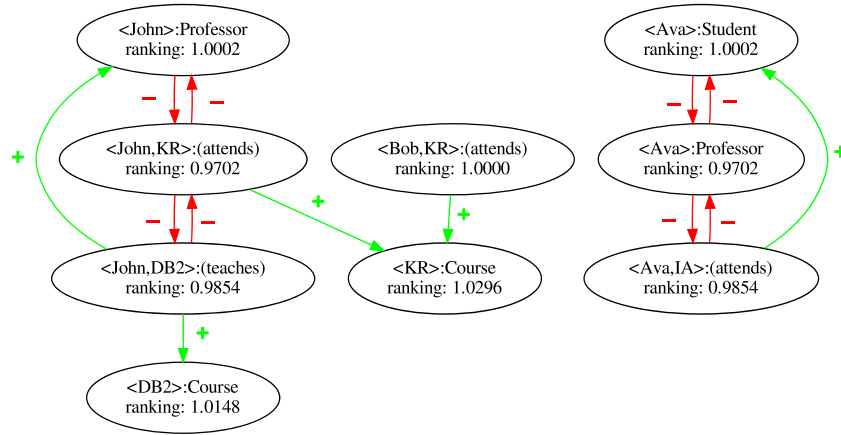


Fig. 2. New ABox assessment after adding Ava : Professor.

become possible, including brave semantics [BR13], ABox Repair (AR) and Intersection ABox Repair (IAR) semantics [LLR⁺10]. These semantics can be enriched by taking into account notions of cardinality [LB16], preference [BBG14], or probability [GM12,CLP18]. Some works [DQ15,TBB⁺17] in the Description Logics community have already investigated the problems of finding the best repairs according to some criteria, and of extracting consistent information that best complies with specific needs. Sik Chun Lam et al. have proposed logic-based methods for repairing TBox axioms of unsatisfiable ontologies [LPSV06a], as well as numerical methods for rewriting and ranking problematic axioms [LPSV06b].

Ranking the nodes of some sort of knowledge graph in terms of interest, quality, or preference is a recurrent problem in many disciplines. Solving these problems requires to somehow quantify the nodes and their interactions. A quantitative approach, rather than a qualitative one, has been used in argumentation frameworks [AB13,BDKM16,BCPR12,Rad18,HT17], belief revision [SSF16], the Web [Kle99,PBMW99], and social networks [dKD08,TND10].

4 Theoretical Framework

Throughout this paper, we will assume that TBoxes are satisfiable. That is, for every TBox \mathcal{T} , the knowledge base $\langle \mathcal{T}, \emptyset \rangle$ is consistent. If a knowledge base $\langle \mathcal{T}, \mathcal{A} \rangle$ is inconsistent, we will assume that the inconsistency is caused by one or more assertions in \mathcal{A} . When humans are faced with such an inconsistent knowledge base $\langle \mathcal{T}, \mathcal{A} \rangle$, they may not be able to quickly pinpoint the “wrong” assertion(s) in \mathcal{A} , because the inconsistency may only become apparent after an involved reasoning process that uses many axioms and assertions. We present

a method for ranking assertions such that higher-ranked assertions are more “truthful” than lower-ranked assertions. Significantly, the ranking only requires some superficial input from the user, who is modeled by the notion of a *credibility function*, which is a mapping from \mathcal{A} to \mathbb{R} . Such a credibility function will be combined with TBox assertions to result in the desired ranking. All definitions that follow are relative to a fixed knowledge base $\langle \mathcal{T}, \mathcal{A} \rangle$ in some fixed Description Logic.

Refuters and supporters

We define what it means for a set B of assertions to support or refute an assertion α not in B . In our framework, assertions will be higher ranked if they have strong supporters and no strong refuters.

Definition 1. *The following definitions are relative to a knowledge base $\langle \mathcal{T}, \mathcal{A} \rangle$, and an assertion $\alpha \in \mathcal{A}$.*

A refuter of α is an inclusion-minimal subset B of \mathcal{A} with the property that $\langle \mathcal{T}, B \rangle$ is consistent and $\langle \mathcal{T}, B \cup \{\alpha\} \rangle$ is inconsistent.

A supporter of α is an inclusion-minimal subset B of \mathcal{A} with the property that $\langle \mathcal{T}, B \rangle$ is consistent, $\alpha \notin B$, and $\langle \mathcal{T}, B \rangle \models \alpha$.

Example 1. Let $\mathcal{T} = \{A \sqsubseteq C, A \sqsubseteq \neg D\}$. Let $\mathcal{A} = \{a : A, a : D, a : C\}$. Then,

- $\{a : A\}$ is a refuter of $a : D$, and a supporter of $a : C$.
- $\{a : C\}$ is neither a refuter nor a supporter of $a : A$.

Proposition 1. *Let $\langle \mathcal{T}, \mathcal{A} \rangle$ be a knowledge base in a monotonic Description Logic, and let $\alpha \in \mathcal{A}$. Then,*

1. *distinct refuters of α are not comparable by set inclusion;*
2. *distinct supporters of α are not comparable by set inclusion; and*
3. *if B is a refuter of α and B' is a supporter of α , then B and B' are not comparable by set inclusion.*

Proof. The proof of the first two items is straightforward. We next prove the last item. Assume B is a refuter of α , and B' is a supporter of α . Consequently,

- (a) $\langle \mathcal{T}, B \rangle$ is consistent;
- (b) $\langle \mathcal{T}, B \cup \{\alpha\} \rangle$ is inconsistent;
- (c) $\langle \mathcal{T}, B' \rangle$ is consistent; and
- (d) $\langle \mathcal{T}, B' \rangle \models \alpha$.

From (c) and (d), it follows $\langle \mathcal{T}, B' \cup \{\alpha\} \rangle$ is consistent.

We first show $B \not\subseteq B'$. Assume for a contradiction that $B \subseteq B'$, hence $B \cup \{\alpha\} \subseteq B' \cup \{\alpha\}$. From (b) and our assumption that the underlying Description Logic is monotonic, it follows that $\langle \mathcal{T}, B' \cup \{\alpha\} \rangle$ is inconsistent, a contradiction.

Finally, we show $B' \not\subseteq B$. Assume for a contradiction $B' \subseteq B$. From (d) and our assumption that the underlying Description Logic is monotonic, it follows $\langle \mathcal{T}, B \rangle \models \alpha$. By (a), it follows $\langle \mathcal{T}, B \cup \{\alpha\} \rangle$ is consistent, which contradicts (b). \square

From here on, we will assume that all Description Logics considered are monotonic.

Aggregated credibility

As explained in the beginning of Section 4, we assume a *credibility function* mapping every assertion α in \mathcal{A} to a real number that can be thought of as the truthfulness associated by an expert to the assertion α .

We will assume that the credibility function induces a function f from $2^{\mathcal{A}}$ to \mathbb{R} by means of some aggregation operator \oplus . Examples of \oplus are min, max, avg, sum. In the theoretical development, we will abstract away from the aggregation operator \oplus and treat f as a basic construct. To keep our framework flexible and general, we do not impose any restrictions on the choice of the credibility function, which can be instantiated and adapted later on to fit a particular setting, for example, as a probability distribution.

ABox assessment

An *ABox assessment* for an ABox \mathcal{A} is a total function $\nu : \mathcal{A} \rightarrow \mathbb{R}$. Informally, our goal is to define ν such that if two ABox assertions are logically conflicting, then the assertion with the higher ν -value is preferred.

Although credibility functions and ABox assessments are both mappings from \mathcal{A} to \mathbb{R} , it is important to understand that they are conceptually distinct. In particular, ABox assessments improve credibility functions by taking into account the axioms of the TBox, via the notions of refuter and supporter defined previously.

Informally, we want that for every $\alpha \in \mathcal{A}$, its value under ν is greater if α has strong supporters, and smaller if α has strong refuters. Here, the strength of a supporter or refuter is recursively determined by the aggregated ν -values of its elements. This can be expressed as follows, where \sim denotes some proportionality that will be made explicit later on:

$$\nu(\alpha) \sim \sum_{\substack{B \text{ is a} \\ \text{supporter} \\ \text{of } \alpha}} \left(f(B) * \sum_{\beta \in B} \nu(\beta) \right) - \sum_{\substack{B \text{ is a} \\ \text{refuter} \\ \text{of } \alpha}} \left(f(B) * \sum_{\beta \in B} \nu(\beta) \right). \quad (1)$$

The expression at the righthand of \sim will be abbreviated $\Sigma^{\pm}(\alpha)$. Informally, $\Sigma^{\pm}(\alpha)$ sums over all supporters and refuters of α . The formulas for supporters and refuters are signed positively and negatively respectively. The magnitude of each supporter's or refuter's contribution is proportional to its aggregated credibility and ABox assessment values. It is significant that (1) is recursive, in the sense that ν appears at both the lefthand and righthand of \sim .

To shorten notations, we define mappings $T : \mathcal{A} \rightarrow 2^{\mathcal{A}}$ and $F : \mathcal{A} \rightarrow 2^{\mathcal{A}}$, as follows:

$$\begin{aligned} T(\alpha) &:= \{B \subseteq \mathcal{A} \mid B \text{ is a supporter of } \alpha\}; \\ F(\alpha) &:= \{B \subseteq \mathcal{A} \mid B \text{ is a refuter of } \alpha\}. \end{aligned}$$

Informally, one can think of T and F as True and False respectively. $B \in T(\alpha)$ is shorthand for “ B is a supporter of α ,” and $B \in F(\alpha)$ is shorthand for “ B is

a refuter of α ". Obviously, the complexity of computing T and F will depend on the underlying Description Logic.

To have a more compact form for $\Sigma^\pm(\alpha)$, we define an indicator function I from $\{T, F\} \times 2^{\mathcal{A}} \times \mathcal{A} \times \mathcal{A}$ to $\{0, 1\}$:

$$I(L, B, \alpha, \beta) = \begin{cases} 1 & \text{if } \beta \in B \text{ and } B \in L(\alpha) \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

Thus, $I(T, B, \alpha, \beta)$ is one if B is a supporter of α that contains β , and is zero otherwise. Likewise, $I(F, B, \alpha, \beta)$ is one if B is a refuter of α that contains β , and is zero otherwise. Note incidentally that $\alpha = \beta$ implies $I(L, B, \alpha, \beta) = 0$, because α does not belong to a supporter or a refuter of itself. Then, for $\alpha \in \mathcal{A}$,

$$\Sigma^\pm(\alpha) = \sum_{\beta \in \mathcal{A}} \nu(\beta) * \left(\sum_{B \subseteq \mathcal{A}} f(B) * (I(T, B, \alpha, \beta) - I(F, B, \alpha, \beta)) \right). \quad (3)$$

Ranking of ABox assertions

An ABox assessment ν induces a ranking of an ABox, as follows: α is ranked higher than β if $\nu(\alpha) > \nu(\beta)$; and α is ranked equal to β if $\nu(\alpha) = \nu(\beta)$. Two ABox assessments ν_1 and ν_2 are *rank-equivalent*, denoted $\nu_1 \overset{\bullet}{\sim} \nu_2$, if they rank assertions in the same way, that is, if for all $\alpha, \beta \in \mathcal{A}$, $\nu_1(\alpha) > \nu_1(\beta)$ if and only if $\nu_2(\alpha) > \nu_2(\beta)$. It is easily verified that if $\nu_1 \overset{\bullet}{\sim} \nu_2$, then $\nu_1(\alpha) = \nu_1(\beta)$ implies $\nu_2(\alpha) = \nu_2(\beta)$. Obviously, $\overset{\bullet}{\sim}$ is an equivalence relation on the set of all ABox assessments. In our study, we will be primarily interested in the set of ABox assessments modulo $\overset{\bullet}{\sim}$. That is, we are not so much interested in the real numbers in the range of two different ABox assessments, but rather in their induced rankings.

5 Framework Instantiation

We will assume $\mathcal{A} = \{\alpha_1, \dots, \alpha_N\}$. In this section, we elaborate an instantiation of the framework in Section 4. With the previous abbreviations, Eq. (1) reads as $\nu(\alpha_i) \sim \Sigma^\pm(\alpha_i)$. We will instantiate \sim as a linear proportionality, i.e., for some numbers a, b, c with $a \neq 0$:

$$\begin{cases} a * \nu(\alpha_1) = b * \Sigma^\pm(\alpha_1) + c \\ a * \nu(\alpha_2) = b * \Sigma^\pm(\alpha_2) + c \\ \vdots \\ a * \nu(\alpha_N) = b * \Sigma^\pm(\alpha_N) + c \end{cases} \quad (4)$$

A more natural formulation may introduce only two numbers d, e and impose $\nu(\alpha_i) = d * \Sigma^\pm(\alpha_i) + e$ for $1 \leq i \leq N$. This is tantamount to fixing $a = 1$. The choice for using three numbers was made for reasons of flexibility, but is not fundamental.

Since we think of (1) as a positive proportionality, we require that a and b have the same sign, say both positive. Let \mathbb{A}^f be the square $N \times N$ matrix such that for $1 \leq i, j \leq N$,

$$\mathbb{A}_{i,j}^f := \sum_{B \subseteq \mathcal{A}} f(B) * (I(T, B, \alpha_i, \alpha_j) - I(F, B, \alpha_i, \alpha_j)). \quad (5)$$

It is easily verified that for every $i \in \{1, \dots, N\}$, $\mathbb{A}_{i,i}^f = 0$.

For $i \in \{1, \dots, N\}$, we introduce a variable x_i for the unknown $\nu(\alpha_i)$. We define $X = [x_1 \ x_2 \ \dots \ x_N]^\top$. Then, using (3), the system of equations (4) can be equivalently written as follows:

$$(a \cdot \mathbb{1} - b \cdot \mathbb{A}^f) \cdot X = [c \ \dots \ c]^\top, \quad (6)$$

where $\mathbb{1}$ is the square identity matrix. Equation (6) has a unique solution if and only if the matrix $a \cdot \mathbb{1} - b \cdot \mathbb{A}^f$ (which does not depend on c) is non-singular (i.e., is invertible). In that case, the solution ABox assessment is given by:

$$X = (a \cdot \mathbb{1} - b \cdot \mathbb{A}^f)^{-1} \cdot [c \ c \ \dots \ c]^\top \quad (7)$$

We should pick $c \neq 0$ to avoid the all-zero solution. Indeed, if $a \cdot \mathbb{1} - b \cdot \mathbb{A}^f$ is non-singular and $c = 0$, then the solution to (7) yields the ABox assessment ν such that $\nu(\alpha_1) = \nu(\alpha_2) = \dots = \nu(\alpha_N) = 0$, which obviously is of no interest for ranking ABox assertions.

Then, it is easily verified that for fixed values of a and b , different choices for c lead to rank-equivalent solutions. Therefore, we can choose $c = 1$ without loss of generality, in which case in the solution, x_i is the sum of all elements in the i th row of $(a \cdot \mathbb{1} - b \cdot \mathbb{A}^f)^{-1}$. Two significant questions remain:

1. For which values of a and b is the square matrix $a \cdot \mathbb{1} - b \cdot \mathbb{A}^f$ non-singular, such that (7) gives us a unique ABox assessment? Obviously, $a \cdot \mathbb{1} - b \cdot \mathbb{A}^f$ is invertible if and only if $\mathbb{1} - \frac{b}{a} \cdot \mathbb{A}^f$ is invertible. Therefore, only the ratio between a and b matters in our subsequent analysis.
2. Are ABox assessments obtained for different values of a, b rank-equivalent?

Graph-theoretical view

From a graph-theoretical viewpoint, \mathbb{A}^f can be interpreted as an edge-weighted directed graph whose vertices are the assertions of the ABox. This view is used in Figures 1 and 2. An edge from α_j to α_i ($i \neq j$) has weight $\mathbb{A}_{i,j}^f$. The task is then to assign a real number to each vertex such that the resulting graph satisfies the sort of equilibrium expressed by (4). The ratio between a and b determines how strongly a vertex is impacted (supported or refuted) by its incoming edges. The question arising is whether an equilibrium (4) can be attained for some values of a, b . It is equally important to examine the robustness of such an equilibrium (if it exists) with respect to rank-equivalence. Ideally, different ABox assessments obtained by small parameter changes should be close modulo rank-equivalence.

6 Solving the Instantiated Framework

In Section 5, we presented the matrix equation (6) whose unique solution, if it exists, yields an ABox assessment. This matrix equation has parameters a, b . A remaining problem is how to pick appropriate values for these parameters. One solution could be to choose values for a, b in an empirical fashion, relying, for instance, on information obtained from some learning experience. Although we have conducted some experiments (cf. Figures 1 and 2), a thorough empirical or experimental validation is left for future research. Instead, we will focus more on theoretical aspects in the current paper. In Section 6.1, we will show a parameterization scheme for a, b that guarantees the existence of a solution for (6): pick $a \gg b > 0$. What is probably more important is that when the ratio a/b grows, the ABox assessments obtained for different a, b become all rank-equivalent. Informally, the parameterization scheme converges to a unique ABox assessment modulo rank-equivalence.

6.1 Solution Existence

Note that all diagonal elements in $a \cdot \mathbb{1} - b \cdot \mathbb{A}^f$ are equal to a . The invertibility of $a \cdot \mathbb{1} - b \cdot \mathbb{A}^f$ can be guaranteed by (some form of) diagonal dominance, a concept that is easily grasped and has turned out to be useful in applications [Var76]. An example is the Levy-Desplanques theorem recalled next.

Theorem 1 (Levy-Desplanques). *A square $N \times N$ matrix A is invertible if for every $i \in \{1, \dots, N\}$, $|a_{ii}| > \sum_{\substack{j=1 \\ j \neq i}}^N |a_{ij}|$.*

From Theorem 1, it immediately follows that (7) has a unique solution if for all $i \in \{1, \dots, N\}$, we have $|a| > |b| * \sum_{j=1}^N |\mathbb{A}_{i,j}^f|$. This is illustrated by Example 2.

Example 2. Assume four assertions $\alpha_1, \alpha_2, \alpha_3, \alpha_4$ such that α_1 and α_2 support one another, while α_3 and α_4 refute one another. Take $b = 1$. If we assume $f(\{\alpha_i\}) = 1$ for $1 \leq i \leq 4$, we obtain:

$$\mathbb{A}^f = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 \\ 0 & 0 & -1 & 0 \end{bmatrix}.$$

Then,

$$a \cdot \mathbb{1} - b \cdot \mathbb{A}^f = \begin{bmatrix} a & -1 & 0 & 0 \\ -1 & a & 0 & 0 \\ 0 & 0 & a & 1 \\ 0 & 0 & 1 & a \end{bmatrix}.$$

It follows from Theorem 1 that this matrix is invertible for $a > 1$, giving the following solution:

$$X = \begin{bmatrix} \frac{1}{a-1} \\ \frac{1}{a-1} \\ \frac{1}{a+1} \\ \frac{1}{a+1} \end{bmatrix}.$$

Although distinct values for a lead to different ABox assessments, it is significant to observe that $\frac{1}{a-1} > \frac{1}{a+1}$ ($a > 1$), and therefore all these ABox assessments ν are rank-equivalent: $\nu(\alpha_1) = \nu(\alpha_2) > \nu(\alpha_3) = \nu(\alpha_4)$. It is intuitively correct that the two assertions that mutually attack one another lose credibility. \square

The invertibility of $a \cdot \mathbb{1} - b \cdot \mathbb{A}^f$ may also follow from Banach's Lemma, which is recalled next.

Theorem 2 (Banach's Lemma). *Let M be a square matrix with the property that $\|M\| < 1$ for some operator norm $\|\cdot\|$. Then the matrix $\mathbb{1} - M$ is invertible and $(\mathbb{1} - M)^{-1} = \mathbb{1} + M + M^2 + M^3 + \dots$.*

Therefore, if we fix $a = 1$ and pick b sufficiently small (which is analogous to fixing $b = 1$ and picking a sufficiently large), then the matrix $\mathbb{1} - b \cdot \mathbb{A}^f$ is invertible, and $(\mathbb{1} - b \cdot \mathbb{A}^f)^{-1} = \mathbb{1} + \mathbb{A}^f + \dots$. If we limit ourselves to the two most significant terms in this expansion, we get $X = (\mathbb{1} + \mathbb{A}^f) \cdot [1 \ 1 \ \dots \ 1]^\top$ as an approximate solution for (7). The ranking obtained by this solution is intuitively meaningful, because it ranks α_i based on the sum of the elements in the i th row of \mathbb{A}^f ; in this row, supporters have a positive sign, and refuters a negative sign.

6.2 Convergence Towards a Fixed Ranking

So it turns out that (6) has a unique solution with $a, b \geq 0$ whenever the ratio a/b is sufficiently large. A desirable property is that different solutions be rank-equivalent, as was the case in Example 2. We now show that this property indeed holds true beyond some threshold value for a/b .

The notion of rank-equivalence obviously extends to any two vectors of real numbers: two such vectors A, B of the same length N are rank-equivalent if for all $1 \leq i, j \leq N$, $a_i < a_j$ if and only if $b_i < b_j$. The proof of the following theorem is in Appendix A. Since, as argued before, only the ratio between a and b matters in our analysis, we can fix b and let a increase in the following theorem.

Theorem 3 (Convergence). *Let M be an $N \times N$ matrix whose diagonal elements are zero. Let $b, c \in \mathbb{R}$. Consider the following equation with real-number parameter a :*

$$(a * \mathbb{1} - b * M) \cdot X = [c \ \dots \ c]^\top. \quad (8)$$

Then there exists $a^ \in \mathbb{R}$ such that*

- for every $a \geq a^*$, the equation (8) has a unique solution; and*

- for all $a_1, a_2 \geq a^*$, if X_1 and X_2 are solutions to (8) for parameters a_1 and a_2 respectively, then X_1 and X_2 are rank-equivalent.

Moreover, such a bound a^* can be computed in polynomial time in the size of N .

Informally, Theorem 3 tells us that for all choices of b and c , there is a threshold value a^* such that all choices of a satisfying $a \geq a^*$ lead to the same ABox assessment modulo rank-equivalence.

6.3 Computational Complexity

We discuss the computational complexity of solving (6). The main difficulty is the computation of the non-diagonal elements in A^f , which are defined by (5). Given α_i and β_j , this requires finding every $B \subseteq \mathcal{A}$ such that $\beta_j \in B$ and B is a supporter or refuter of α_i . In general, the number of supporters or refuters can be exponential in the size of \mathcal{A} , because an ABox of size $2N$ can have $\binom{2N}{N}$ supporters or refuters not comparable by set inclusion. However, in practice, a fixed upper bound on the size of supporters or refuters may be implied by the TBox, in a way captured by the following definition.

Definition 2. *We say that a TBox \mathcal{T} is conflict-bounded if there exists a polynomial-time computable positive integer m such that for every knowledge-base $\langle \mathcal{A}, \mathcal{T} \rangle$ and $\alpha \in \mathcal{A}$, if B is a refuter or supporter of α , then $|B| \leq m$.*

Conflict-boundedness arises in practice. For example, for every DL-Lite TBox, every supporter or refuter has cardinality ≤ 1 because each conflict in DL-Lite involves at most two assertions. Conflict-boundedness is also guaranteed in $\mathcal{EL}_{\perp nr}$ (\mathcal{EL}_{\perp} with non-recursive empty concepts) defined in [Ros11].

Theorem 4. *Let \mathcal{T} be a conflict-bounded TBox such that ABox consistency with respect to \mathcal{T} can be checked in polynomial time. Then, for every ABox \mathcal{A} , the matrix A^f can be computed in polynomial time (in the size of \mathcal{A}) if f is polynomial-time computable.*

Proof. Let $N := |\mathcal{A}|$. Recall that A^f is an $N \times N$ matrix given by (5). Since \mathcal{T} is conflict-bounded, say with bound m , for every $B \subseteq \mathcal{A}$, if $|B| > m$, then $I(T, B, \alpha_i, \alpha_j) = I(F, B, \alpha_i, \alpha_j) = 0$. Therefore, the sum in (5) must only consider subsets $B \subseteq \mathcal{A}$ with $|B| \leq m$, which are polynomially many and polynomial-time computable. Furthermore, since ABox consistency checking is in polynomial time by the hypothesis of the theorem, it can be tested in polynomial time whether any such B containing α_j is a supporter or a refuter of α_i (or neither of both). The computation of $f(B)$ is also in polynomial time by the hypothesis of the theorem. It is now obvious that any entry of A^f can be computed in polynomial time. \square

An inspection of the preceding proof reveals that Theorem 4 remains valid if, throughout its statement, we replace polynomial time by some higher complexity upper bound (e.g., polynomial space or exponential time).

7 Summary and Future Work

In this work, we have proposed a novel framework to rank the assertions of an inconsistent ABox in terms of their degree of truthfulness. This ranking takes into account an expert’s credibility function as well as the logical axioms of the TBox. The theoretical framework can work with any Description Logic, but to gain computational tractability, restrictions have to be imposed.

The framework is implementable using existing libraries for matrix operations and description logics. In the future, we aim to extend our prototype in order to empirically evaluate our theory.

A next step is to use our framework for repairing database and knowledge-base systems. Our ranking of ABox assertions can be extended in several ways to a ranking of repairs, using some notion of Pareto optimality. For example, if two facts α and β contradict one another and α is ranked higher than β , then a repair containing α is preferred over a repair containing β (all other facts being equal). This brings us in the realm of prioritized database repairing [Wij19], which we see as a promising way to enrich existing DL repair semantics (like IAR and AR). By narrowing the search space of possible repairs, we may also hope to gain efficiency compared to existing repair semantics.

Acknowledgments We thank the anonymous reviewers for their helpful comments.

References

- [AB13] Leila Amgoud and Jonathan Ben-Naim. Ranking-based semantics for argumentation frameworks. In Weiru Liu, V. S. Subrahmanian, and Jef Wijsen, editors, *Scalable Uncertainty Management - 7th International Conference, SUM 2013, Washington, DC, USA, September 16-18, 2013. Proceedings*, volume 8078 of *Lecture Notes in Computer Science*, pages 134–147. Springer, 2013.
- [BBG14] Meghyn Bienvenu, Camille Bourgaux, and François Goasdoué. Querying inconsistent description logic knowledge bases under preferred repair semantics. In Carla E. Brodley and Peter Stone, editors, *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, July 27 -31, 2014, Québec City, Québec, Canada*, pages 996–1002. AAAI Press, 2014.
- [BCPR12] Richard Booth, Martin Caminada, Mikolaj Podlaskowski, and Iyad Rahwan. Quantifying disagreement in argument-based reasoning. In Wiebe van der Hoek, Lin Padgham, Vincent Conitzer, and Michael Winikoff, editors, *International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2012, Valencia, Spain, June 4-8, 2012 (3 Volumes)*, pages 493–500. IFAAMAS, 2012.
- [BDKM16] Elise Bonzon, Jérôme Delobelle, Sébastien Konieczny, and Nicolas Maudet. A comparative study of ranking-based semantics for abstract argumentation. In Dale Schuurmans and Michael P. Wellman, editors, *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, pages 914–920. AAAI Press, 2016.

- [BR13] Meghyn Bienvenu and Riccardo Rosati. Tractable approximations of consistent query answering for robust ontology-based data access. In Francesca Rossi, editor, *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013*, pages 775–781. IJCAI/AAAI, 2013.
- [CLP18] Marco Calautti, Leonid Libkin, and Andreas Pieris. An operational approach to consistent query answering. In Jan Van den Bussche and Marcelo Arenas, editors, *Proceedings of the 37th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, Houston, TX, USA, June 10-15, 2018*, pages 239–251. ACM, 2018.
- [dKD08] Cristobald de Kerchove and Paul Van Dooren. The pagetrust algorithm: How to rank web pages when negative links are allowed? In *Proceedings of the SIAM International Conference on Data Mining, SDM 2008, April 24-26, 2008, Atlanta, Georgia, USA*, pages 346–352. SIAM, 2008.
- [DQ15] Jianfeng Du and Guilin Qi. Tractable computation of representative ABox repairs in description logic ontologies. In Songmao Zhang, Martin Wirsing, and Zili Zhang, editors, *Knowledge Science, Engineering and Management - 8th International Conference, KSEM 2015, Chongqing, China, October 28-30, 2015, Proceedings*, volume 9403 of *Lecture Notes in Computer Science*, pages 28–39. Springer, 2015.
- [GM12] Sergio Greco and Cristian Molinaro. Probabilistic query answering over inconsistent databases. *Ann. Math. Artif. Intell.*, 64(2-3):185–207, 2012.
- [HT17] Anthony Hunter and Matthias Thimm. Probabilistic reasoning with abstract argumentation frameworks. *J. Artif. Intell. Res.*, 59:565–611, 2017.
- [Kle99] Jon M. Kleinberg. Authoritative sources in a hyperlinked environment. *J. ACM*, 46(5):604–632, 1999.
- [LB16] Andrei Lopatenko and Leopoldo E. Bertossi. Complexity of consistent query answering in databases under cardinality-based and incremental repair semantics (extended version). *CoRR*, abs/1605.07159, 2016.
- [LLR⁺10] Domenico Lembo, Maurizio Lenzerini, Riccardo Rosati, Marco Ruzzi, and Domenico Fabio Savo. Inconsistency-tolerant semantics for description logics. In Pascal Hitzler and Thomas Lukasiewicz, editors, *Web Reasoning and Rule Systems - Fourth International Conference, RR 2010, Bressanone/Brixen, Italy, September 22-24, 2010. Proceedings*, volume 6333 of *Lecture Notes in Computer Science*, pages 103–117. Springer, 2010.
- [LPSV06a] Sik Chun Lam, Jeff Z. Pan, Derek H. Sleeman, and Wamberto Weber Vasconcelos. A fine-grained approach to resolving unsatisfiable ontologies. In *2006 IEEE / WIC / ACM International Conference on Web Intelligence (WI 2006), 18-22 December 2006, Hong Kong, China*, pages 428–434. IEEE Computer Society, 2006.
- [LPSV06b] Sik Chun Lam, Jeff Z. Pan, Derek H. Sleeman, and Wamberto Weber Vasconcelos. Ontology inconsistency handling : Ranking and rewriting axioms. Technical report, University of Aberdeen, 2006.
- [PBMW99] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The pagerank citation ranking: Bringing order to the web. Technical Report 1999-66, Stanford InfoLab, November 1999. Previous number = SIDL-WP-1999-0120.
- [Rad18] Badran Raddaoui. On the measure of conflicts: an argumentation-based framework. *J. Appl. Non Class. Logics*, 28(2-3):240–259, 2018.

- [Ros11] Riccardo Rosati. On the complexity of dealing with inconsistency in description logic ontologies. In Toby Walsh, editor, *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, pages 1057–1062. IJCAI/AAAI, 2011.
- [SSF16] Gerardo I. Simari, Paulo Shakarian, and Marcelo A. Falappa. A quantitative approach to belief revision in structured probabilistic argumentation. *Ann. Math. Artif. Intell.*, 76(3-4):375–408, 2016.
- [TBB⁺17] Abdelmoutia Telli, Salem Benferhat, Mustapha Bourahla, Zied Bouraoui, and Karim Tabia. Polynomial algorithms for computing a single preferred assertional-based repair. *Künstliche Intell.*, 31(1):15–30, 2017.
- [TND10] Vincent A. Traag, Yurii E. Nesterov, and Paul Van Dooren. Exponential ranking: Taking into account negative links. In Leonard Bolc, Marek Makowski, and Adam Wierzbicki, editors, *Social Informatics - Second International Conference, SocInfo 2010, Laxenburg, Austria, October 27-29, 2010. Proceedings*, volume 6430 of *Lecture Notes in Computer Science*, pages 192–202. Springer, 2010.
- [Var76] Richard S. Varga. On recurring theorems on diagonal dominance. *Linear Algebra and its Applications*, 13(1):1 – 9, 1976.
- [Wij19] Jef Wijsen. Foundations of query answering on inconsistent databases. *SIGMOD Rec.*, 48(3):6–16, 2019.

A Proof of Theorem 3

Proof (of Theorem 3). Clearly, if x_c is the solution to

$$(a * \mathbb{1} - b * M) \cdot X = c * [1 \ 1 \ \dots \ 1]^\top$$

with $c > 0$, and x_1 is the solution to

$$(a * \mathbb{1} - b * M) \cdot X = [1 \ 1 \ \dots \ 1]^\top,$$

then $x_c = c * x_1$. Since x_c and x_1 are rank-equivalent, we can fix $c = 1$ without loss of generality. Define a' as follows:

$$a' := 1 + |b| * (N - 1) \left(\max_{1 \leq i, j \leq N} |M_{ij}| \right).$$

By the Levy-Desplanques theorem, for every $a \geq a'$, the matrix $a * \mathbb{1} - b * M$ is invertible. For every $a \geq a'$, we write x^a for the unique solution of the equation

$$(a * \mathbb{1} - b * M) \cdot X = [1 \ 1 \ \dots \ 1]^\top.$$

For $a \geq a'$, define $\delta_{ij}(a) := x_i^a - x_j^a$, where x_i^a is the i th coordinate of x^a .

We will show that there is a value $a^* \geq a'$ such that for all $a_1, a_2 \geq a^*$, for all $1 \leq i < j \leq N$, $\delta_{ij}(a_1)$ and $\delta_{ij}(a_2)$ have the same sign, which implies that x^{a_1} and x^{a_2} are rank-equivalent.

Define the following matrix $S_{|i}$ in function of a , for $1 \leq i \leq N$:

$$(S_{|i}(a))_{\ell k} = \begin{cases} (a * \mathbb{1} - b * M)_{\ell k} & \text{if } k \neq i \\ 1 & \text{if } k = i \end{cases} \quad (9)$$

That is, $S_{|i}$ is obtained from $a * \mathbb{1} - b * M$ by replacing the i th column with $[1 \ 1 \ \dots \ 1]^\top$. By Cramer's Rule, we have

$$x_i^a = \frac{\det(S_{|i}(a))}{\det(a * \mathbb{1} - b * M)}, \quad (10)$$

and consequently

$$\delta_{ij}(a) = \frac{\det(S_{|i}(a)) - \det(S_{|j}(a))}{\det(a * \mathbb{1} - b * M)}. \quad (11)$$

Since the determinants $\det(\cdot)$ are polynomial expressions, the following are all polynomials of degree at most N :

- $p_i(a) = \det(S_{|i}(a))$;
- $p_j(a) = \det(S_{|j}(a))$; and
- $p(a) = \det(a * \mathbb{1} - b * M)$.

If $a \geq a'$, then since the polynomial $p(a)$ does not vanish, the sign of $\delta_{ij}(a)$ is fully determined by the expression

$$\det(S_{|i}(a)) - \det(S_{|j}(a)). \quad (12)$$

Let $p_{ij} = p_i - p_j$ for $1 \leq i < j \leq N$. We determine a^* such that for all $a \geq a^*$ and for all $1 \leq i < j \leq N$, the sign of $p_{ij}(a)$ does not change (i.e., is either always positive or always negative). Note that the sign of p_{ij} not changing is equivalent to the sign of p_{ji} not changing, so we can assume $i < j$ without loss of generality.

In the remainder of the proof, we show that the desired a^* exists and can be computed in polynomial time in N . Pick $N+1$ real numbers $V = \{a_1, \dots, a_{N+1}\}$, all greater than a' . Compute

$$p_{ij}(a_k) = \det(S_{|i}(a_k)) - \det(S_{|j}(a_k)) \quad (13)$$

for all $a_k \in V$ and $1 \leq i < j \leq N$. The set

$$\{p_{ij}(a_k) \mid 1 \leq i < j \leq N, 1 \leq k \leq N+1\}$$

can be computed in polynomial time in N , because it involves $N(N+1)$ determinants (i.e., $\det(S_{|i}(a_k))$ for $1 \leq i \leq N$ and $1 \leq k \leq N+1$), each of which can be computed in polynomial time in N . For all $1 \leq i < j \leq N$, the polynomial p_{ij} can be computed from $\{(a_1, p_{ij}(a_1)), \dots, (a_{N+1}, p_{ij}(a_{N+1}))\}$ in polynomial time using, for example, Lagrange interpolation. The number of polynomials to compute is $\frac{N(N-1)}{2}$ (i.e., polynomially many), and each of them has at most

N coefficients. We will represent the polynomial p_{ij} by its coefficients, i.e., by $\langle (p_{ij})_0, \dots, (p_{ij})_{n_{ij}} \rangle$ where each $(p_{ij})_\ell$ is the coefficient of degree ℓ , and $n_{ij} \leq N$ is the polynomial's degree. We now define a_{ij}^* as

$$a_{ij}^* := \max \left(a', 2 + \max_{0 \leq \ell \leq n_{ij}-1} \frac{-(p_{ij})_\ell}{|(p_{ij})_{n_{ij}}|} \right).$$

By Cauchy's bound on positive real roots of polynomials, if x_0 is a root of p_{ij} , then $x_0 < a_{ij}^*$. This implies that if $a_1, a_2 > a_{ij}^*$, then $p_{ij}(a_1)$ and $p_{ij}(a_2)$ have the same sign (i.e., either both positive or both negative), and therefore, by definition of δ_{ij} , we have $x_i^{a_1} < x_j^{a_1}$ if and only if $x_i^{a_2} < x_j^{a_2}$. Finally, let

$$a^* = \max_{1 \leq i < j \leq N} a_{ij}^*,$$

which can be computed in polynomial time. By our construction, it follows that for all $a_1, a_2 > a^*$, the solutions x^{a_1}, x^{a_2} exist and are rank-equivalent.

Finally, we incidentally note that a slightly better bound is obtained by letting

$$a^* = \max \left(a', 2 + \max_{1 \leq i < j \leq N, 0 \leq \ell \leq n_{ij}-1} \frac{-(p_{ij})_\ell}{|(p_{ij})_{n_{ij}}|} \right). \quad (14)$$

This concludes the proof. \square