The Thirty-Third International
FLAIRS Conference (FLAIRS-33)

On Dealing with Conflicting, Uncertain and Partially Ordered Ontologies

Sihem Belabbes, Salem Benferhat
CRIL, Univ. Artois & CNRS, Lens, France
{belabbes, benferhat}@cril.fr

Abstract

We focus on handling conflicting and uncertain information in lightweight ontologies, where uncertainty is represented in a possibilistic logic setting. We use DL-Lite, a tractable fragment of Description Logic, to specify terminological knowledge (i.e., TBox). We assume the TBox to be stable and coherent, while its combination with a set of assertional facts (i.e., ABox) may be inconsistent. We address the problem of dealing with conflicts when the reliability relation between sources is only partially ordered. We propose to represent the uncertain ABox as a symbolic weighted base, where a strict partial order is applied on the weights. In this context, we provide a strategy for computing a single repair for the ABox, called the partial possibilistic repair. The idea is to consider all compatible bases of a partially preordered ABox (which intuitively encode total extensions of a partial order), compute their associated possibilistic repairs, before intersecting those repairs to produce a single consistent sub-base.

Introduction

Description Logics (DLs) are a family of logic-based knowledge representation languages successfully used in a broad range of applications, thanks to the good trade-off they offer between expressive power and computational complexity. In particular, DLs are widely used for representing and reasoning about ontologies. Inconsistency management in formal ontologies, especially those specified in DL-Lite, a lightweight fragment of description logics, has been tackled from different angles. One line of research supports reasoning in the presence of incomplete, uncertain, qualitative and prioritized information using possibility theory (Dubois, Prade, and Schockaert 2017; Finger et al. 2017).

For instance, fuzzy extensions have been proposed for DLs (Borgwardt and Peñaloza 2017; Bobillo and Straccia 2018; Straccia 2013) and for DL-Lite (Pan et al. 2007; Straccia 2006). There has also been possibilistic extensions of DLs (Dubois, Mengin, and Prade 2006; Qi et al. 2011) as well as probabilistic ones (Baader et al. 2019; Borgwardt, Ceylan, and Lukasiewicz 2018; Lutz and Schröder 2010).

Furthermore, a framework for possibilistic DL-Lite has been proposed (Benferhat and Bouraoui 2017). It draws inspiration from Standard Possibilistic Logic (Dubois and Prade 2015) in which inconsistent and uncertain information is prioritized by way of a total preorder. Basically, weights in the unit interval $[0, 1]$ (seen as an ordinal scale) are attached to ABox assertions to represent different reliability levels of the information. The higher the weight, the more certain or reliable is the assertion. A weight (or degree) is considered as a lower bound on the assertion's certainty (or priority) level. A nice feature of possibilistic DL-Lite is that query answering is tractable despite the fact that the expressiveness of standard DL-Lite is enhanced with a total preorder over the assertions.

Nevertheless, in applications such as ontologies, information is typically obtained from multiple sources having conflicting opinions. This implies applying a partial order instead of a total order over the weights assigned to assertions. Note that the order relation applied to weights is a strict partial order, i.e., there are no ties between weights. However the order relation on the corresponding assertions is a partial preorder, since the same weight could be attached to more than one assertion (i.e., ties between assertions are allowed). Hence the corresponding ABox is partially preordered.

Extensions of Standard Possibilistic Logic have been proposed to support reasoning with partially preordered information, mainly using the notion of compatible bases. In (Benferhat, Lagrue, and Papini 2004), possibilistic inference is revisited by assigning symbolic weights defined over a partially ordered uncertainty scale to propositional logic formulas. This idea has also been explored in (Benferhat, Dubois, and Prade 1995; Touazi, Cayrol, and Dubois 2015). However, computational complexity is expensive (Δ_p^2-hard), making such approaches not suitable in a context where queries need to be answered efficiently.

A natural question is whether standard possibilistic DL-Lite (Benferhat and Bouraoui 2017) can be extended to take into account partially preordered knowledge, without increasing computational complexity. Recently, an efficient method, called "Elect", has been proposed for the case where the ABox is partially preordered (Belabbes, Benferhat, and Chomicki 2019). In essence, Elect computes a single repair for an inconsistent ABox and does so in polynomial time. It has been shown that Elect generalises the well-known Intersection ABox Repair (IAR) semantics (Lembo et al. 2010) defined for a flat ABox and the so-called non-defeated repair (Benferhat, Bouraoui, and Tabia 2015; Benferhat, Dubois, and Prade 1998) defined for a totally preordered ABox. Basically, a partially preordered ABox is interpreted as a family of totally preordered ABoxes such that a single repair can be computed for each one of them. The intersection of all those repairs produces a single repair for the partially preordered ABox.

In this paper, we follow an approach similar to the Elect method and propose a possibilistic strategy for computing a single repair for the ABox, called the partial possibilistic repair. To achieve this aim, we consider a family of compatible ABoxes (which amount to a family of possibilistic DL-Lite ABoxes) and compute the possibilistic repair associated with each one of them. We then intersect those possibilistic repairs to obtain a single repair for the partially preordered weighted ABox.

We start the paper by briefly recalling the basics of DL-Lite in description logic followed by its extension to possibilistic logic. We introduce our method for computing a possibilistic repair for a partially preordered weighted ABox. We then conclude and discuss future work.

The Description Logic DL-Lite

Description Logics (DLs) are a family of knowledge representation languages that have been successfully applied in various domains and specifically in formalising ontologies. The lightweight fragments of DLs such as DL-Lite (Calvanese et al. 2007) are particularly popular, since they offer a good trade-off between expressive power and computational complexity. For instance, query answering from a DL-Lite knowledge base can be carried out efficiently using query rewriting (Kontchakov et al. 2010). In this paper, we present the DL-Lite$_R$ dialect of DL-Lite, without loss of generality.

A Knowledge Base (KB) is built upon three finite and mutually disjoint sets C, R and I, which contain respectively *concept names*, *role names* and *individual names*. The DL-Lite$_R$ language is defined according to the following rules:

$$R \longrightarrow P \mid P^- \qquad E \longrightarrow R \mid \neg R$$
$$B \longrightarrow A \mid \exists R \qquad C \longrightarrow B \mid \neg B$$

where A denotes a concept name, P is a role name, and P^- is the *inverse* of P. Also, R stands for a *basic role* and E denotes a *complex role*. Furthermore, B is a *basic concept* while C is a *complex concept*.

A DL-Lite KB $\mathcal{K}$ is a tuple $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$, where:

- $\mathcal{T}$ is a finite set of *inclusion axioms*, also known as TBox. An inclusion axiom on concepts (resp. on roles) is a statement of the form $B \sqsubseteq C$ (resp. $R \sqsubseteq E$). Concept inclusions are said to be *negative inclusion axioms* if they contain the symbol "¬" to the right of the inclusion, otherwise they are called *positive inclusion axioms*.
- $\mathcal{A}$ is a finite set of *assertions* (ground facts), also known as ABox. An assertion is a statement of the form $A(a)$ or $P(a, b)$, where $a, b \in \mathsf{I}$.

A KB $\mathcal{K}$ is said to be *consistent* if it admits at least one model, otherwise it is *inconsistent*. A TBox $\mathcal{T}$ is *incoherent* if there is a concept name $A \in \mathsf{C}$ such that A is empty in every model of $\mathcal{T}$, otherwise it is *coherent*.
Henceforth, we shall refer to DL-Lite$_R$ simply as DL-Lite.

We shall use the following running example throughout the paper and adapt it as needed.

Example 1 *Let $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ be a DL-Lite KB.*
Let $\mathcal{T} = \{A \sqsubseteq \neg B, B \sqsubseteq \neg C, C \sqsubseteq \neg D\}$ be a TBox.
Let $\mathcal{A} = \{A(a),\ A(b),\ B(a),\ B(c),\ C(a),\ C(b),\ D(a),\ D(b),\ D(c),\ E(a)\}$ be a flat ABox (i.e., no weights are assigned to assertions).
One can easily check that $\mathcal{K}$ is inconsistent. For instance, individual 'a' belongs to both concepts A and B. This contradicts the negative axiom $A \sqsubseteq \neg B$.

□

There exist various strategies in the literature, called inconsistency-tolerant semantics, for reasoning with inconsistent KBs (e.g. (Baget et al. 2016; Calvanese et al. 2010; Bienvenu and Bourgaux 2016; Trivela, Stoilos, and Vassalos 2019)). Basically, they proceed by computing a single or several consistent sub-bases of the ABox, known as repairs, then use those repairs to perform reasoning tasks such as query answering. The most well-known strategies are the ABox Repair (AR) semantics and the Intersection ABox Repair (IAR) semantics (Lembo et al. 2010). In AR, queries are evaluated separately over all the repairs, then the sets of answers are intersected. Thus a query answer is considered valid if it can be entailed from every repair of the ABox. In IAR, queries are evaluated over one consistent sub-base of the ABox obtained from the intersection of all the repairs. Other strategies have been proposed such as the so-called non-defeated semantics (Benferhat, Bouraoui, and Tabia 2015), which amounts to a prioritized version of the IAR semantics.

In the present paper, we focus on possibilistic repairs, especially in the case of partially preordered knowledge. Note that a repair is usually defined as an inclusion-maximal subset of the ABox that is consistent with respect to the TBox. Here we use the term repair in the possibilistic context even for a subset of assertions that is not maximal, provided it is consistent with respect to the TBox.

In the next section, we recall the underpinnings of standard possibilistic DL-Lite.

Possibilistic DL-Lite Knowledge Base

Possibilistic Description Logics (Hollunder 1995; Dubois, Mengin, and Prade 2006) are extensions of standard Description Logics frameworks based on possibility theory that support reasoning with uncertain and inconsistent knowledge. Extensions to possibilistic DL-Lite (Benferhat and Bouraoui 2017) have recently been proposed for the lightweight fragments DL-Lite. The main idea consists in assigning priority degrees (or weights) to TBox axioms and ABox assertions to express their relative certainty (or confidence) in an inconsistent KB. The inconsistency degree of the KB can then be computed from those weights, making provision for possibilistic inference.

In this section, we consider a possibilistic DL-Lite KB $\mathcal{WK} = \langle \mathcal{T}, \mathcal{WA} \rangle$, henceforth referred to as weighted KB. We assume axioms in $\mathcal{T}$ to be fully certain (or fully reliable) while assertions in $\mathcal{WA}$ (for weighted ABox) are equipped with priority degrees from the interval $]0, 1]$ as follows:

$$\mathcal{WA} = \{(f, \alpha) : f \text{ is a DL-Lite assertion}, \alpha \in]0, 1]\}.$$

Assertions in $\mathcal{WA}$ with priority degree $\alpha = 1$ are fully certain and cannot be questioned, whereas assertions with priority degree $0 < \alpha < 1$ are somewhat certain. Assertions with higher degrees are more certain than those with lower degrees. We ignore assertions whose degree $\alpha = 0$.

Henceforth, for any given weighted ABox $\mathcal{B}$, we shall denote by $\mathcal{B}^*$ the set of assertions without priority degrees.

We also assume that the weighted KB $\mathcal{WK}$ may be inconsistent. Furthermore we assume the TBox component to be coherent and stable, thus the inconsistency of $\mathcal{WK}$ is caused by conflicts between assertions of $\mathcal{WA}$ w.r.t. axioms of $\mathcal{T}$.

An assertional conflict is defined as an inclusion-minimal subset of assertions that is inconsistent with the TBox, where inconsistency is understood in the sense of standard DL-Lite. Formally:

Definition 1 *Let $\mathcal{WK} = \langle \mathcal{T}, \mathcal{WA} \rangle$ be a weighted KB. A sub-base $\mathcal{C} \subseteq \mathcal{WA}$ is an assertional conflict in $\mathcal{WK}$ iff:*

- *$\langle \mathcal{T}, \mathcal{C}^* \rangle$ is inconsistent, and*
- *$\forall f \in \mathcal{C}^*$, $\langle \mathcal{T}, \mathcal{C}^* \setminus \{f\} \rangle$ is consistent.*

Let $\mathcal{C}(\mathcal{WA})$ denote the set of all assertional conflicts of $\mathcal{WA}$. It is important to highlight that computing the set of conflicts is done in polynomial time in DL-Lite (Calvanese et al. 2010). Furthermore, assertional conflicts in coherent DL-Lite KBs are binary, i.e., $\forall \mathcal{C} \in \mathcal{C}(\mathcal{WA}), |\mathcal{C}| = 2$ (Calvanese et al. 2010). Thus we denote a conflict by a pair $\mathcal{C}_{ij} = \{(f_i, \alpha_i), (f_j, \alpha_j)\}$, where $(f_i, \alpha_i), (f_j, \alpha_j) \in \mathcal{WA}$, and say that assertions $f_i, f_j \in \mathcal{WA}^*$ are conflicting w.r.t. $\mathcal{T}$.

Example 2 *We equip the ABox of Example 1 with weights. Let $\mathcal{WK} = \langle \mathcal{T}, \mathcal{WA} \rangle$ be a weighted KB, where the TBox $\mathcal{T} = \{A \sqsubseteq \neg B, B \sqsubseteq \neg C, C \sqsubseteq \neg D\}$, and the weighted ABox*

$$\mathcal{WA} = \left\{ \begin{array}{l} (A(a), .9), (A(b), .9), (B(c), .8), \\ (E(a), .7), (D(b), .6), (C(a), .5), \\ (D(a), .4), (B(a), .3), (D(c), .3), \\ (C(b), .1) \end{array} \right\}$$

The set of assertional conflicts of $\mathcal{WA}$ is given by:

$$\mathcal{C}(\mathcal{WA}) = \left\{ \begin{array}{l} \{(A(a), .9), (B(a), .3)\}, \\ \{(D(b), .6), (C(b), .1)\}, \\ \{(C(a), .5), (D(a), .4)\}, \\ \{(C(a), .5), (B(a), .3)\} \end{array} \right\}$$

□

As shall be made clear later, we are interested in the highest priority degree where inconsistency is met in the ABox, known as the inconsistency degree. Formally:

Definition 2 *Let $\mathcal{WK} = \langle \mathcal{T}, \mathcal{WA} \rangle$ be a weighted KB. Consider a weight $\beta \in]0, 1]$. We denote by:*

- *$\mathcal{A}^{\geq \beta} = \{f : (f, \alpha) \in \mathcal{WA}, \alpha \geq \beta\}$ the β-cut of $\mathcal{WA}$.*
- *$\mathcal{A}^{> \beta} = \{f : (f, \alpha) \in \mathcal{WA}, \alpha > \beta\}$ the strict β-cut of $\mathcal{WA}$.*

The inconsistency degree of $\mathcal{WA}$, denoted by $Inc(\mathcal{WA})$, is:

$$Inc(\mathcal{WA}) = \begin{cases} 0 & \textit{iff } \langle \mathcal{T}, \mathcal{WA}^* \rangle \textit{ is consistent} \\ \beta & \textit{iff } \langle \mathcal{T}, \mathcal{A}^{\geq \beta} \rangle \textit{ is inconsistent} \\ & \textit{and } \langle \mathcal{T}, \mathcal{A}^{> \beta} \rangle \textit{ is consistent} \end{cases}$$

We illustrate this notion on our running example.

Example 3 *One can easily check that for $\beta = 0.4$, we have:*

- *$\mathcal{A}^{> \beta} = \{A(a), A(b), B(c), E(a), D(b), C(a)\}$ is consistent w.r.t. $\mathcal{T}$, whereas*
- *$\mathcal{A}^{\geq \beta} = \mathcal{A}^{> \beta} \cup \{D(a)\}$ is inconsistent w.r.t. $\mathcal{T}$.*

Therefore: $Inc(\mathcal{WA}) = 0.4$.

□

The inconsistency degree serves as a means for restoring consistency of the ABox. This is due to the fact that only assertions with a certainty degree that is strictly higher than the inconsistency degree are included in the possibilistic repair, which ensures safety of the results. Moreover, this has the advantage of being efficient. Indeed, for a weighted ABox $\mathcal{WA}$, $Inc(\mathcal{WA})$ can be computed tractably using $log_2(n)$ (where n is the number of different weights in $\mathcal{WA}$) consistency checks of a classical ABox (without weights).

The possibilistic repair, henceforth referred to as π-repair, is formally defined as follows:

Definition 3 *Let $\mathcal{WK} = \langle \mathcal{T}, \mathcal{WA} \rangle$ be a weighted KB and $Inc(\mathcal{WA})$ the inconsistency degree. The π-repair of $\mathcal{WA}$, denoted by $\pi(\mathcal{WA})$, is given by:*

$$\pi(\mathcal{WA}) = \{f : (f, \alpha) \in \mathcal{WA}, \alpha > Inc(\mathcal{WA})\}.$$

The π-repair $\pi(\mathcal{WA})$ is composed of those assertions of $\mathcal{WA}$ of which the priority degree is strictly higher than $Inc(\mathcal{WA})$. Hence by Definition 2, $\pi(\mathcal{WA})$ is consistent with $\mathcal{T}$. Also note that priority degrees are omitted in $\pi(\mathcal{WA})$. Moreover, when $\mathcal{WK}$ is consistent (i.e., $Inc(\mathcal{WA}) = 0$), then $\pi(\mathcal{WA})$ amounts to $\mathcal{WA}^*$ (i.e., the ABox without priority degrees).

Example 4 *The π-repair of $\mathcal{WA}$ is:*
$\pi(\mathcal{WA}) = \{A(a), A(b), B(c), C(a), D(b), E(a)\}$.

□

So far, we have considered weighted ABoxes such that the weights attached to assertions can be used to induce a total preorder on the ABox. Next, we scale the results to the case where priority degrees are partially ordered.

Partially Preordered Knowledge Base

In this section, we still assume TBox axioms are fully reliable. However, priorities associated with ABox assertions are partially preordered, i.e., reliability levels associated with some assertions may be incomparable. This is often the case when information is obtained from multiple sources. Thus we may not be able to decide on a preference between two assertions f_i and f_j because according to one source, assertion f_i should be preferred to f_j, whereas according to another source, it should be the opposite.

Let us define a partially ordered uncertainty scale $\mathbb{L} = (U, \rhd)$, over a non-empty set of elements $U = \{u_1, \ldots, u_n\}$, called a partially ordered set (POS), and a strict partial order $\rhd$ (irreflexive and transitive relation).

Intuitively, elements of U represent symbolic priority degrees applied to assertions. We assume that U contains a special element $\mathbb{1}$ representing full certainty, such that for all $u_i \in U$, $\mathbb{1} \rhd u_i$. Moreover, if $u_i \not\rhd u_j$ and $u_j \not\rhd u_i$, we say that u_i and u_j are incomparable and write $u_i \sim u_j$.

A partially preordered DL-Lite KB is a triple $\mathcal{K}_\rhd = \langle \mathcal{T}, \mathcal{A}_\rhd, \mathbb{L} \rangle$, where $\mathcal{A}_\rhd = \{(f_i, u_i) : f_i$ is a DL-Lite assertion, $u_i \in U\}$ and $\mathbb{L} = (U, \rhd)$.
Given two assertions $(f_i, u_i), (f_j, u_j) \in \mathcal{A}_\rhd$, we shall sometimes abuse notations and write $f_i \rhd f_j$ to mean $u_i \rhd u_j$ and $f_i \sim f_j$ to mean $u_i \sim u_j$.

Compatible Bases

A natural way of representing a partially preordered ABox is to consider the set of all compatible ABoxes, i.e., those that preserve the strict preference ordering between assertions, in the spirit of proposals made in the context of propositional logic (Benferhat, Lagrue, and Papini 2004). Formally:

Definition 4 *Let $\mathbb{L} = (U, \rhd)$ be a partially ordered uncertainty scale. Let $\mathcal{K}_\rhd = \langle \mathcal{T}, \mathcal{A}_\rhd, \mathbb{L} \rangle$ be a partially preordered DL-Lite KB. Let $\mathcal{WK} = \langle \mathcal{T}, \mathcal{WA} \rangle$ be a weighted KB, obtained from $\mathcal{K}_\rhd$ by replacing each element u by a real number in the interval $]0, 1]$, where:*

$$\mathcal{WA} = \{(f, \alpha) : (f, u) \in \mathcal{A}_\rhd, \alpha \in]0, 1]\}.$$

The weighted ABox $\mathcal{WA}$ is said to be compatible with $\mathcal{A}_\rhd$ iff: $\forall (f_i, \alpha_i), (f_j, \alpha_j) \in \mathcal{WA}$, if $f_i \rhd f_j$ then $\alpha_i > \alpha_j$.

Note that compatible bases are not unique, actually there is an infinite number thereof. In fact, the actual values of weights do not really matter, only the ordering between assertions matters, as shall be shown later.

Example 5 *Let $\mathbb{L} = (U, \rhd)$ be a partially ordered uncertainty scale defined over the set $U = \{u_1, \ldots, u_4\}$, such that: $u_4 \rhd u_3 \rhd u_1$, $u_4 \rhd u_2 \rhd u_1$ and $u_2 \sim u_3$.*
Let $\mathcal{K}_\rhd = \langle \mathcal{T}, \mathcal{A}_\rhd, \mathbb{L} \rangle$ be a partially preordered KB.
Let $\mathcal{T} = \{A \sqsubseteq \neg B, B \sqsubseteq \neg C, C \sqsubseteq \neg D\}$. Let the ABox be:

$$\mathcal{A}_\rhd = \left\{ \begin{array}{l} (A(a), u_4), (A(b), u_4), (B(c), u_4), (C(a), u_3), \\ (D(b), u_3), (E(a), u_3), (C(b), u_2), (B(a), u_1), \\ (D(a), u_1), (D(c), u_1) \end{array} \right\}$$

Consider the set of weights $\{.2, .4, .6, .8\}$. The following bases are compatible with $\mathcal{A}_\rhd$:

$$\mathcal{WA}_1 = \left\{ \begin{array}{l} (A(a), .8), (A(b), .8), (B(c), .8), \\ (C(a), .6), (D(b), .6), (E(a), .6), \\ (C(b), .4), \\ (B(a), .2), (D(a), .2), (D(c), .2) \end{array} \right\}$$

$$\mathcal{WA}_2 = \left\{ \begin{array}{l} (A(a), .8), (A(b), .8), (B(c), .8), \\ (C(b), .6), \\ (C(a), .4), (D(b), .4), (E(a), .4), \\ (B(a), .2), (D(a), .2), (D(c), .2) \end{array} \right\}$$

$$\mathcal{WA}_3 = \left\{ \begin{array}{l} (A(a), .8), (A(b), .8), (B(c), .8), \\ (C(a), .6), (D(b), .6), (E(a), .6), (C(b), .6), \\ (B(a), .4), (D(a), .4), (D(c), .4) \end{array} \right\}$$

□

Computing the Partial Possibilistic Repair

We are interested in computing a single repair for a partially preordered ABox. However, the family of compatible ABoxes is infinite, which means that selecting one compatible ABox over others would be arbitrary. A better approach for computing the partial possibilistic repair consists in:

(i) defining the compatible ABoxes (Definition 4) with weights defined over $]0, 1]$,

(ii) computing the π-repair associated with each compatible ABox (Definition 3),

(iii) and finally intersecting all π-repairs.

This ensures the safety of the results since all compatible ABoxes are taken into account.

Definition 5 *Let $\mathbb{L} = (U, \rhd)$ be a partially ordered uncertainty scale. Let $\mathcal{K}_\rhd = \langle \mathcal{T}, \mathcal{A}_\rhd, \mathbb{L} \rangle$ be a partially preordered DL-Lite KB. Let $\mathcal{F}(\mathcal{A}_\rhd) = \{\pi(\mathcal{WA}) : \mathcal{WA}$ is compatible with $\mathcal{A}_\rhd\}$ be the set of π-repairs associated with all compatible bases of $\mathcal{A}_\rhd$ (given by Definition 3).*
The partial possibilistic repair, denoted by $\pi(\mathcal{A}_\rhd)$, is:

$$\pi(\mathcal{A}_\rhd) = \bigcap \{\pi(\mathcal{WA}) : \pi(\mathcal{WA}) \in \mathcal{F}(\mathcal{A}_\rhd)\}.$$

Namely, $\pi(\mathcal{A}_\rhd) = \{f : (f, u) \in \mathcal{A}_\rhd, \forall \mathcal{WA}$ compatible with $\mathcal{A}_\rhd, f \in \pi(\mathcal{WA})\}$.

Note that weights are omitted in the partial possibilistic repair $\pi(\mathcal{A}_\rhd)$, similarly to the π-repair $\pi(\mathcal{WA})$.

The set $\mathcal{F}(\mathcal{A}_\rhd)$ is infinite because there are infinitely many weighted ABoxes that are compatible with the partially preordered ABox $\mathcal{A}_\rhd$. However, we do not need to consider all compatible bases of $\mathcal{A}_\rhd$ in order to compute the partial possibilistic repair $\pi(\mathcal{A}_\rhd)$. Indeed, it is enough to consider only the compatible bases (and their associated repairs) that define a different ordering between assertions. This is captured by the following lemma.

Lemma 1 *Let $\mathcal{WA}_1$ be a weighted ABox. Let $S = \{\alpha : (f, \alpha) \in \mathcal{WA}_1\}$ be the set of weights attached to assertions of $\mathcal{WA}_1$. Consider an assignment function $\omega : S \longrightarrow]0, 1]$ s.t. $\forall \alpha_1, \alpha_2 \in S, \alpha_1 \geq \alpha_2$ iff $\omega(\alpha_1) \geq \omega(\alpha_2)$.*
Let $\mathcal{WA}_2 = \{(f, \omega(\alpha)) : (f, \alpha) \in \mathcal{WA}_1\}$ be a weighted ABox obtained by applying assignment function ω to the weights attached to assertions of $\mathcal{WA}_1$. Then:

$$\pi(\mathcal{WA}_1) = \pi(\mathcal{WA}_2).$$

In Lemma 1, although $\mathcal{WA}_2$ is different from $\mathcal{WA}_1$, the former ABox preserves the ordering on the latter's assertions. Thus $\mathcal{WA}_2$ is said to be order-preserving and in this case, the two weighted bases generate the same repairs.

Proof: Let us show that $Inc(\mathcal{WA}_1) = \beta$ iff $Inc(\mathcal{WA}_2) = \omega(\beta)$. First note that if $\mathcal{C}_{12} = \{(f_1, \alpha_1), (f_2, \alpha_2)\}$ and $\mathcal{C}_{34} = \{(f_3, \alpha_3), (f_4, \alpha_4)\}$ are two conflicts of $\mathcal{WA}_1$, then obviously $\mathcal{C}'_{12} = \{(f_1, \omega(\alpha_1)), (f_2, \omega(\alpha_2))\}$ and $\mathcal{C}'_{34} = \{(f_3, \omega(\alpha_3)), (f_4, \omega(\alpha_4))\}$ are also two conflicts of $\mathcal{WA}_2$. Then, by definition of the function $\omega(\cdot)$, if we have $min\{\alpha : (f, \alpha) \in \mathcal{C}_{12}\} = \alpha_1$ (resp. α_2), then we also have $min\{\omega(\alpha) : (f, \omega(\alpha)) \in \mathcal{C}'_{12}\} = \omega(\alpha_1)$ (resp. $\omega(\alpha_2)$).

Similarly, if $min\{\alpha : (f, \alpha) \in \mathcal{C}_{12}\} > min\{\alpha : (f, \alpha) \in \mathcal{C}_{34}\}$, then $min\{\omega(\alpha) : (f, \omega(\alpha)) \in \mathcal{C}'_{12}\} > min\{\omega(\alpha) : (f, \omega(\alpha)) \in \mathcal{C}'_{34}\}$. Thus if $Inc(\mathcal{WA}_1) = \beta$, then trivially $Inc(\mathcal{WA}_2) = \omega(\beta)$.

Assume $Inc(\mathcal{WA}_1) = \beta$. Let $(f, \alpha) \in \mathcal{WA}_1$ s.t. $\alpha > \beta$. Then $f \in \pi(\mathcal{WA}_1)$. By definition of $\omega(\cdot)$, we get $\omega(\alpha) > \omega(\beta) = Inc(\mathcal{WA}_2)$. This means $f \in \pi(\mathcal{WA}_2)$. Similarly, let $(f, \alpha) \in \mathcal{WA}_1$ s.t. $\alpha \leq \beta$. Then $f \notin \pi(\mathcal{WA}_1)$. Again by definition of $\omega(\cdot)$, we get $\omega(\alpha) \leq \omega(\beta) = Inc(\mathcal{WA}_2)$. This means $f \notin \pi(\mathcal{WA}_2)$.

Therefore we conclude that $\pi(\mathcal{WA}_1) = \pi(\mathcal{WA}_2)$. ■

Let us illustrate these notions on our running example.

Example 6 *Thanks to Lemma 1, in order to compute the repair $\pi(\mathcal{A}_\triangleright)$, it is enough to consider only the three bases $\mathcal{WA}_1$, $\mathcal{WA}_2$ and $\mathcal{WA}_3$ as compatible bases of $\mathcal{A}_\triangleright$. Their associated π-repairs are given by:*

- $\pi(\mathcal{WA}_1) = \{A(a), A(b), B(c), C(a), D(b), E(a)\}$.
- $\pi(\mathcal{WA}_2) = \{A(a), A(b), B(c), C(b)\}$.
- $\pi(\mathcal{WA}_3) = \{A(a), A(b), B(c)\}$.

The partial possibilistic repair is:

$$\pi(\mathcal{A}_\triangleright) = \bigcap_{i=1\ldots3} \pi(\mathcal{WA}_i) = \{A(a), A(b), B(c)\}.$$

□

We conclude that reasoning (i.e., answering queries) from a partially preordered inconsistent KB amounts to replacing the original ABox $\mathcal{A}_\triangleright$ with its repair $\pi(\mathcal{A}_\triangleright)$.

Discussion of a Characterization

In previous work (Belabbes, Benferhat, and Chomicki 2019), a method called Elect has been introduced to compute a single repair for an inconsistent DL-Lite KB, where a partial preorder is directly applied to the assertions of the ABox, i.e., no weights are attached to them. The idea consists in viewing a partially preordered ABox as a family of totally preordered ABoxes to which the so-called non-defeated semantics (Benferhat, Bouraoui, and Tabia 2015) is applied. This produces non-defeated repairs which are then intersected in order to produce a single repair for the initial partially preordered ABox. An equivalent characterization has been provided in order to avoid eliciting all total preorders. It relies on the notion of elected assertion (Belabbes, Benferhat, and Chomicki 2019), which is defined as an assertion that is strictly preferred (according to the partial preorder) to all its opponents (i.e., assertions involved in a conflict with it). It has been shown that the set of all elected assertions corresponds indeed to the Elect repair, that it is consistent with the TBox, and that its computation can be achieved in polynomial time in DL-Lite.

The question now is whether a similar approach can be followed in a possibilistic setting. In other words, we would like to define an equivalent characterization for computing the partial possibilistic repair $\pi(\mathcal{A}_\triangleright)$, without enumerating all compatible bases $\mathcal{WA}_i$ of $\mathcal{A}_\triangleright$. Due to space limitations, the details are left for future work and we provide here a brief discussion.

Recall that partially ordered symbolic weights are applied to the assertions of $\mathcal{A}_\triangleright$ and that numeric weights in the unit interval $]0, 1]$ are applied to the assertions of the compatible bases $\mathcal{WA}_i$. The former case can be viewed as a partial preorder applied to the assertions (without weights) and the latter case can be viewed as a total preorder applied to the assertions. It follows that the idea of the Elect method can indeed be used. However we argue that the notion of accepted assertion in a possibilistic setting should be weaker than the notion of elected assertion. Basically, an assertion of a partially preordered possibilistic ABox $\mathcal{A}_\triangleright$ is said to be accepted if it is strictly preferred to at least one assertion of each assertional conflict of $\mathcal{A}_\triangleright$. Clearly, if an assertion is accepted in the possibilistic setting, then it is also elected.

We illustrate this notion on our running example where the set of assertional conflicts of $\mathcal{A}_\triangleright$ is:

$$\mathcal{C}(\mathcal{A}_\triangleright) = \left\{ \begin{array}{l} \{(A(a), u_4), (B(a), u_1)\}, \\ \{(C(a), u_3), (D(a), u_1)\}, \\ \{(D(b), u_3), (C(b), u_2)\}, \\ \{(C(a), u_3), (B(a), u_1)\} \end{array} \right\}$$

It is easy to see that assertions $(A(a), u_4)$, $(A(b), u_4)$ and $(B(c), u_4)$ are strictly preferred to at least one assertion of each conflict, since the symbolic weight u_4 is strictly preferred to all other weights. Hence these assertions are all accepted. This corresponds to the result of Example 6, i.e., $\pi(\mathcal{A}_\triangleright)$, where weights are omitted.

Note that the base $\pi(\mathcal{A}_\triangleright)$ is obviously consistent w.r.t. the TBox. Indeed, since the π-repair $\pi(\mathcal{WA}_i)$ of each compatible base $\mathcal{WA}_i$ of $\mathcal{A}_\triangleright$ is consistent, the intersection of all π-repairs is necessarily consistent.

In addition, by construction of $\pi(\mathcal{A}_\triangleright)$, it is straightforward to see that when the partial order $\triangleright$ is a total order denoted by $\succ$, then $\pi(\mathcal{A}_\triangleright)$ collapses with the π-repair $\pi(\mathcal{A}_\succ)$.

Conclusion

In this paper, we proposed an extension of possibilistic DL-Lite to partially preordered inconsistent knowledge bases. Basically, a partially preordered ABox is interpreted as a family of compatible weighted ABoxes for which possibilistic repairs are computed then intersected to produce a single repair for the partially preordered ABox. Due to space limitations, the characterization of the partial possibilistic repair and its computational properties are left for future work.

In the future, we also plan to enhance the productivity of the partial repair by considering the closure of possibilistic repairs associated with the compatible ABoxes. A crucial question is whether the computation of the closed partial possibilistic repair can be achieved in polynomial time in DL-Lite. An idea consists in reducing the problem to answering an instance checking query. We plan to investigate whether polynomial methods for computing repairs in the flat and prioritized cases are also polynomial with a partial order for DLs in general (more expressive than DL-Lite).

Moreover we plan to develop an efficient tool for query answering from ontologies representing Southeast Asian traditional dances. Dance videos are semantically enriched by domain experts through annotations w.r.t. the ontology (i.e., the TBox). Conflicts may emerge when the same video is annotated differently by several experts. Experts may assign confidence degrees to their annotations which corresponds to a totally preordered ABox. However different experts may not share the same meaning of confidence scales which corresponds to applying a partial preorder to the ABox.

Acknowledgement This work is supported by the European project H2020-MSCA-RISE AniAge: High dimensional heterogeneous data based animation techniques for Southeast Asian intangible cultural heritage digital content.

References

Baader, F.; Ecke, A.; Kern-Isberner, G.; and Wilhelm, M. 2019. The complexity of the consistency problem in the probabilistic description logic $\mathcal{ALC}^{\mathsf{me}}$. In *FroCoS, London, UK*, 167–184.

Baget, J.; Benferhat, S.; Bouraoui, Z.; Croitoru, M.; Mugnier, M.; Papini, O.; Rocher, S.; and Tabia, K. 2016. A general modifier-based framework for inconsistency-tolerant query answering. In *KR, Cape Town, South Africa*, 513–516.

Belabbes, S.; Benferhat, S.; and Chomicki, J. 2019. Elect: An inconsistency handling approach for partially preordered lightweight ontologies. In *LPNMR, Philadelphia, USA*, 210–223.

Benferhat, S., and Bouraoui, Z. 2017. Min-based possibilistic DL-Lite. *Journal of Logic and Computation* 27(1):261–297.

Benferhat, S.; Bouraoui, Z.; and Tabia, K. 2015. How to select one preferred assertional-based repair from inconsistent and prioritized DL-Lite knowledge bases? In *IJCAI, Buenos Aires, Argentina*, 1450–1456.

Benferhat, S.; Dubois, D.; and Prade, H. 1995. How to infer from inconsistent beliefs without revising? In *International Joint Conference on Artificial Intelligence*, 1449–1457. Morgan Kaufmann.

Benferhat, S.; Dubois, D.; and Prade, H. 1998. *Some syntactic approaches to the handling of inconsistent knowledge bases : A comparative study. Part 2 : the prioritized case*, volume 24. Physica-Verlag, Heidelberg. 473–511.

Benferhat, S.; Lagrue, S.; and Papini, O. 2004. Reasoning with partially ordered information in a possibilistic logic framework. *Fuzzy Sets and Systems* 144(1):25–41.

Bienvenu, M., and Bourgaux, C. 2016. Inconsistency-tolerant querying of description logic knowledge bases. In *Reasoning Web: Logical Foundation of Knowledge Graph Construction and Query Answering*, volume 9885, 156–202. LNCS. Springer.

Bobillo, F., and Straccia, U. 2018. Reasoning within fuzzy OWL 2 EL revisited. *Fuzzy Sets and Systems* 351:1–40.

Borgwardt, S., and Peñaloza, R. 2017. Fuzzy description logics - a survey. In *Scalable Uncertainty Management (SUM)*, 31–45.

Borgwardt, S.; Ceylan, İ. İ.; and Lukasiewicz, T. 2018. Recent advances in querying probabilistic knowledge bases. In *IJCAI, Stockholm, Sweden*, 5420–5426.

Calvanese, D.; De Giacomo, G.; Lembo, D.; Lenzerini, M.; and Rosati, R. 2007. Tractable reasoning and efficient query answering in description logics: The DL-Lite family. *Journal of Automated Reasoning* 39(3):385–429.

Calvanese, D.; Kharlamov, E.; Nutt, W.; and Zheleznyakov, D. 2010. Evolution of DL-Lite knowledge bases. In *International Semantic Web Conference (1)*, 112–128.

Dubois, D., and Prade, H. 2015. Possibility theory and its applications: Where do we stand? In *Springer Handbook of Computational Intelligence*. 31–60.

Dubois, D.; Mengin, J.; and Prade, H. 2006. Possibilistic uncertainty and fuzzy features in description logic. a preliminary discussion. *Fuzzy Logic and the Semantic Web. Volume 1 of Capturing Intelligence* 101–113.

Dubois, D.; Prade, H.; and Schockaert, S. 2017. Generalized possibilistic logic: Foundations and applications to qualitative reasoning about uncertainty. *Artificial Intelligence* 252:139–174.

Finger, M.; Godo, L.; Prade, H.; and Qi, G. 2017. Advances in weighted logics for artificial intelligence. *International Journal of Approximate Reasoning* 88:385–386.

Hollunder, B. 1995. An alternative proof method for possibilistic logic and its application to terminological logics. *International Journal of Approximate Reasoning* 12(2):85–109.

Kontchakov, R.; Lutz, C.; Toman, D.; Wolter, F.; and Zakharyaschev, M. 2010. The combined approach to query answering in DL-Lite. In *KR, Toronto, Canada*, 247–257.

Lembo, D.; Lenzerini, M.; Rosati, R.; Ruzzi, M.; and Savo, D. F. 2010. Inconsistency-tolerant semantics for description logics. In *International Conference on Web Reasoning and Rule Systems*, volume 6333 of *LNCS*, 103–117.

Lutz, C., and Schröder, L. 2010. Probabilistic description logics for subjective uncertainty. In *KR, Toronto, Canada*.

Pan, J. Z.; Stamou, G. B.; Stoilos, G.; and Thomas, E. 2007. Expressive querying over fuzzy DL-Lite ontologies. In *DL workshop, Bressanone, Italy*.

Qi, G.; Ji, Q.; Pan, J. Z.; and Du, J. 2011. Extending description logics with uncertainty reasoning in possibilistic logic. *International Journal of Intelligent Systems* 26(4):353–381.

Straccia, U. 2006. Towards top-k query answering in description logics: The case of DL-Lite. In *JELIA, Liverpool, UK*, 439–451.

Straccia, U. 2013. *Foundations of Fuzzy Logic and Semantic Web Languages*. Chapman & Hall/CRC.

Touazi, F.; Cayrol, C.; and Dubois, D. 2015. Possibilistic reasoning with partially ordered beliefs. *J. of Applied Logic* 13(4):770–798.

Trivela, D.; Stoilos, G.; and Vassalos, V. 2019. Query rewriting for DL ontologies under the ICAR semantics. In *RuleML+RR, Bolzano, Italy*, 144–158.