Partially Preordered Inconsistent Lightweight Ontologies in Possibility Theory

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

Abstract

In the context of DL-Lite, a family of lightweight fragments of description logics, this paper introduces an extension of standard possibilistic DL-Lite to the case where the knowledge base is partially preordered. We consider the assertional base, a.k.a. ABox, to be represented as a symbolic weighted base and we assume a strict partial preorder is applied on the weights. We introduce a tractable method for computing a single repair for a partially preordered weighted ABox. Basically, this repair is computed from possibilistic repairs associated with compatible bases of a partially preordered ABox, which intuitively encode all possible extensions of a partial preorder. We provide an equivalent characterization using the notion of π-accepted assertions, which ensures that the possibilistic repair is computed in polynomial time.

Introduction

Possibility theory has been widely studied since the seminal work of Zadeh (Zadeh 1978). Basically, it is an uncertainty theory that handles incomplete, uncertain, qualitative and prioritized information and supports reasoning in the presence of inconsistency (Dubois, Prade, and Schockaert 2017; Finger et al. 2017). Possibility theory has strong connections with ordinal conditional functions (Spohn 2014) as well as with consonant belief functions (Fagin et al. 2003; Dempster 1967; Shafer 1976).

Standard Possibilistic Logic (SPL) (Dubois and Prade 2015) provides a natural framework for reasoning with inconsistent and uncertain information that is prioritized by way of a total preorder. In essence, SPL is a weighted logic where formulas are propositional logic formulas, each of which is assigned a weight in the unit interval $[0, 1]$ seen as an ordinal scale. A weight (or degree) is considered as a lower bound on the formula's certainty (or priority) level.

A research domain that has gained considerable interest is that of inconsistency management in formal ontologies, in particular those specified in the lightweight fragments of description logics (DLs) known as DL-Lite. 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). Moreover, possibilistic extensions of DLs (Dubois, Mengin, and Prade 2006; Qi et al. 2011) alongside probabilistic extensions (Baader et al. 2019; Borgwardt, Ceylan, and Lukasiewicz 2018; Lutz and Schröder 2010) have also been proposed.

Recently, a framework for possibilistic DL-Lite has been proposed (Benferhat and Bouraoui 2017). Basically, ABox assertions are assigned weights reflecting the fact that some pieces of information are deemed as more reliable than others. A nice feature of possibilistic DL-Lite is that query answering is tractable. Hence, there is no extra cost despite the fact that the expressiveness of standard DL-Lite is enhanced with a total preorder over ABox assertions.

Nonetheless, in several applications and notably ontologies, reliability is partially defined, usually as a result of obtaining information from multiple sources that do not share the same opinions. This implies the application of a partial preorder instead of a total preorder over the weights assigned to formulas or assertions.

Extensions of SPL have been proposed to support reasoning with partially preordered information. In (Benferhat, Lagrue, and Papini 2004), the core notions of SPL such as possibilistic inference are revisited by assigning degrees to propositional logic formulas that belong to a partially ordered uncertainty scale instead of the unit interval $[0, 1]$. In (Touazi, Cayrol, and Dubois 2015), the idea of assigning partially ordered symbolic weights to beliefs is also studied extensively. The main drawback of such approaches is that their computational complexity is expensive (at least Δ_p^2), which makes them not suitable in a context where queries need to be efficiently answered.

Against this background, we are interested in proposing an extension of standard possibilistic DL-Lite (Benferhat and Bouraoui 2017) to the case where knowledge is partially preordered, without increasing its computational complexity. The idea of handling inconsistency in partially preordered lightweight ontologies has been recently investigated in (Belabbes, Benferhat, and Chomicki 2019). An efficient method, called "Elect", has been proposed for computing a single repair for a partially preordered ABox. Elect offers an extension of the well-known IAR semantics (Lembo et al. 2010) for partially preordered ABoxes. Basically, a partially preordered ABox is interpreted as a family of totally preordered ABoxes for which repairs can be computed. These repairs are then intersected to produce a single repair for the partially preordered ABox.

A natural question is then whether possibilistic DL-Lite can be extended to partial preorders in a tractable way (in the spirit of Elect). This paper provides a positive answer.

To achieve this aim, we consider a family of compatible ABoxes (which amount to possibilistic DL-Lite ABoxes) and compute the possibilistic repair associated with each compatible base. Finally, we compute a single repair for the partially preordered weighted ABox from the intersection of all possibilistic repairs. Our main contribution is the provision of an equivalent characterization that identifies accepted assertions, called π-accepted, without explicitly computing all the compatible bases. We show that the set of π-accepted assertions is consistent and that it can be computed in polynomial time. Moreover, we show that when the preference relation is a total preorder, the produced repair amounts to the possibilistic repair as computed in standard possibilistic DL-Lite.

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

The Description Logic DL-Lite

The Description Logic DL-Lite (Calvanese et al. 2007) is a family of knowledge representation languages that have gained popularity in several application domains such as formalizing lightweight ontologies, thanks to their expressive power and good computational properties. 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 a finite set of *concept names* C, a finite set of *role names* R and a finite set of *individual names* I, where sets C, R and I are mutually disjoint. The DL-Lite$_R$ language is defined according to the following BNF rules:

$R \longrightarrow P \mid P^-$ $\quad E \longrightarrow R \mid \neg R$

$B \longrightarrow A \mid \exists R$ $\quad C \longrightarrow B \mid \neg B$

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 knowledge base $\mathcal{K}$ is composed of two components, $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$, 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 "$\neg$" 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 knowledge base $\mathcal{K}$ 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*.

In the rest of this paper, we shall refer to DL-Lite$_R$ as DL-Lite to simplify notations.

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. □

Several strategies have been proposed for reasoning with inconsistent KBs (e.g. (Baget et al. 2016; Calvanese et al. 2010; Bienvenu and Bourgaux 2016; Trivela, Stoilos, and Vassalos 2019). The main idea consists in computing repairs for the ABox, where a repair is defined as a maximal (w.r.t. set inclusion) subset of the ABox that is consistent with respect to the TBox. Amongst these strategies, one can cite the ABox Repair (AR) semantics (Lembo et al. 2010) in which queries are evaluated over the intersection of all the repairs. There is also the Intersection ABox Repair (IAR) semantics (Lembo et al. 2010) which queries one consistent subbase of the ABox obtained from the intersection of all the repairs. Furthermore, the so-called non-defeated repair semantics (Benferhat, Bouraoui, and Tabia 2015) amounts to a prioritized version of IAR semantics.

In the present paper, we shall focus on possibilistic repairs, especially in the case of partially preordered knowledge. Let us first 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 knowledge base. The inconsistency degree of the knowledge base can then be computed from those weights, making provision for possibilistic inference.

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

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

Assertions in $\mathcal{WA}$ with a priority degree $\alpha = 1$ are considered as fully certain and cannot be questionable, whereas assertions with a priority degree $0 < \alpha < 1$ are somewhat certain. Assertions with higher priority degrees are more certain than those with lower priority degrees. We ignore assertions whose degree $\alpha = 0$, thus only assertions that are somewhat certain are stated explicitly.

In the rest of this paper, for any given weighted assertional base $\mathcal{B}$, we shall denote by $\mathcal{B}^*$ the set of assertions without priority degrees. We shall denote by $\mathcal{WK}^*$ the KB whose ABox component is the set of assertions $\mathcal{B}^*$.

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 a minimal (w.r.t. set inclusion) 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:

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

We denote by $\mathcal{C}(\mathcal{WA})$ the set of all assertional conflicts of $\mathcal{WA}$. It is important to note 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 knowledge bases are binary, i.e., $\forall \mathcal{C} \in \mathcal{C}(\mathcal{WA}), |\mathcal{C}| = 2$ (Calvanese et al. 2010). Thus we denote an assertional conflict by a pair $\mathcal{C} = \{(f_1, \alpha_1), (f_2, \alpha_2)\}$, where $(f_1, \alpha_1), (f_2, \alpha_2) \in \mathcal{WA}$, and say that assertions $f_1, f_2 \in \mathcal{WA}^*$ are conflicting w.r.t. $\mathcal{T}$.

Example 2
We continue Example 1 and equip the ABox 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

$$\mathcal{WA} = \left\{ \begin{array}{l} (A(a), 0.9), (A(b), 0.9), \\ (B(c), 0.8), \\ (E(a), 0.7), \\ (D(b), 0.6), \\ (C(a), 0.5), \\ (D(a), 0.4), \\ (B(a), 0.3), (D(c), 0.3), \\ (C(b), 0.1) \end{array} \right\}$$

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

$$\mathcal{C}(\mathcal{WA}) = \left\{ \begin{array}{l} \{(A(a), 0.9), (B(a), 0.3)\}, \\ \{(D(b), 0.6), (C(b), 0.1)\}, \\ \{(C(a), 0.5), (D(a), 0.4)\}, \\ \{(C(a), 0.5), (B(a), 0.3)\} \end{array} \right\}$$

□

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 $\beta \in]0, 1]$.
Let $\mathcal{A}^{\geq \beta} = \{f \mid (f, \alpha) \in \mathcal{WA}, \alpha \geq \beta\}$ denote the β-cut of the weighted assertional base $\mathcal{WA}$.
Let $\mathcal{A}^{> \beta} = \{f \mid (f, \alpha) \in \mathcal{WA}, \alpha > \beta\}$ denote the strict β-cut of $\mathcal{WA}$.
The inconsistency degree of $\mathcal{WA}$, denoted by $Inc(\mathcal{WA})$, is:

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

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 an inconsistent ABox. This is due to the fact that only assertions whose certainty degree is strictly higher than the inconsistency degree are included in the possibilistic repair, which ensures safety of the results. Moreover, this method has the advantage of being efficient. Indeed, for a weighted ABox $\mathcal{WA}$, $Inc(\mathcal{WA})$ can be computed in a tractable way using $log_2(n)$ (where n is the number of different weights in $|\mathcal{WA}|$) consistency checks of a classic 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:

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

The π-repair $\pi(\mathcal{WA})$ is composed of those assertions of $\mathcal{WA}$ whose 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. In the next section, we scale the results to the case where priority degrees are partially preordered.

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 and g because according to one source, assertion f should be preferred to g, whereas according to another source, it should be the opposite.

Let us introduce the notion of partially ordered uncertainty scale $\mathbb{L} = (U, \rhd)$, defined over a non-empty set of elements U, called a partially ordered set (POS), and a strict partial preorder $\rhd$ (irreflexive and transitive relation).

Intuitively, elements of U represent priority degrees applied to ABox assertions. We assume that U contains a special element denoted by $\mathbb{1}$, where $\mathbb{1}$ represents full certainty, such that for all $u \in U$, $\mathbb{1} \rhd u$. Moreover, if $u \not\rhd v$ and $v \not\rhd u$, we say that u and v are incomparable and denote it by $u \sim v$.

A partially preordered DL-Lite KB is a triple $\mathcal{K}_{\rhd} = \langle \mathcal{T}, \mathcal{A}_{\rhd}, \mathbb{L} \rangle$, where $\mathcal{A}_{\rhd} = \{(f, u) \mid f$ is a DL-Lite assertion, $u \in U\}$ and $\mathbb{L} = (U, \rhd)$.

Given two assertions $(f, u), (g, v) \in \mathcal{A}_{\rhd}$, we shall abuse notations and write $f \rhd g$ to mean $u \rhd v$ and $f \sim g$ to mean $u \sim v$.

Compatible bases

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

Definition 4
Let $\mathbb{L} = (U, \rhd)$ be an 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 unit interval $]0, 1]$, where:

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

The weighted ABox $\mathcal{WA}$ is said to be compatible with $\mathcal{A}_{\rhd}$ if and only if:

$$\forall (f_i, \alpha_i), (f_j, \alpha_j) \in \mathcal{WA}, \textit{ if } f_i \rhd f_j \textit{ then } \alpha_i > \alpha_j.$$

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 an uncertainty scale defined over the set $U = \{u_1, \ldots, u_n\}$, $n \geq 4$, 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 $\mathcal{T} = \{A \sqsubseteq \neg B, B \sqsubseteq \neg C, C \sqsubseteq \neg D\}$.

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

Consider a set of weights $\{\alpha_1, \ldots, \alpha_m\}$, $m \geq 4$, such that: $1 > \alpha_4 > \alpha_3 > \alpha_2 > \alpha_1 > 0$.

The following bases are compatible with $\mathcal{A}_{\rhd}$:

$$\mathcal{WA}_1 = \left\{ \begin{array}{l} (A(a), \alpha_4), (A(b), \alpha_4), (B(c), \alpha_4), \\ (C(a), \alpha_3), (D(b), \alpha_3), (E(a), \alpha_3), \\ (C(b), \alpha_2), \\ (B(a), \alpha_1), (D(a), \alpha_1), (D(c), \alpha_1) \end{array} \right\}$$

$$\mathcal{WA}_2 = \left\{ \begin{array}{l} (A(a), \alpha_4), (A(b), \alpha_4), (B(c), \alpha_4), \\ (C(b), \alpha_3), \\ (C(a), \alpha_2), (D(b), \alpha_2), (E(a), \alpha_2), \\ (B(a), \alpha_1), (D(a), \alpha_1), (D(c), \alpha_1) \end{array} \right\}$$

$$\mathcal{WA}_3 = \left\{ \begin{array}{l} (A(a), \alpha_3), (A(b), \alpha_3), (B(c), \alpha_3), \\ (C(a), \alpha_2), (D(b), \alpha_2), \\ (E(a), \alpha_2), (C(b), \alpha_2), \\ (B(a), \alpha_1), (D(a), \alpha_1), (D(c), \alpha_1) \end{array} \right\}$$

□

Computing partially preordered 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 partially preordered repair consists in:

(i) defining the compatible ABoxes (Definition 4),

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

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

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

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

Note that weights are omitted in the partially preordered 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 partially preordered 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 $S = \{\alpha \mid (f, \alpha) \in \mathcal{WA}_1\}$ be the set of weights attached to assertions of $\mathcal{WA}_1$.
Consider an assignment function $\omega : S \longrightarrow]0, 1]$ such that $\forall \alpha_1, \alpha_2 \in S, \alpha_1 \geq \alpha_2$ iff $\omega(\alpha_1) \geq \omega(\alpha_2)$.
Let $\mathcal{WA}_2 = \{(f, \omega(\alpha)) \mid (f, \alpha) \in \mathcal{WA}_1\}$ be a weighted

ABox obtained by applying the assignment function ω to the weights attached to assertions of $\mathcal{WA}_1$. Then:

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

Although ABox $\mathcal{WA}_2$ is different from ABox $\mathcal{WA}_1$, the former 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:
It is easy to see that $Inc(\mathcal{WA}_1) = \beta$ iff $Inc(\mathcal{WA}_2) = \omega(\beta)$.

Note first that if $\mathcal{C}_1 = \{(f_1, \alpha_1), (f_2, \alpha_2)\}$ and $\mathcal{C}_2 = \{(f_3, \alpha_3), (f_4, \alpha_4)\}$ are two conflicts of $\mathcal{WA}_1$, then obviously $\mathcal{C}'_1 = \{(f_1, \omega(\alpha_1)), (f_2, \omega(\alpha_2))\}$ and $\mathcal{C}'_2 = \{(f_3, \omega(\alpha_3)), (f_4, \omega(\alpha_4))\}$ are also two conflicts of $\mathcal{WA}_2$.

By definition of function $\omega(.)$, if we have $min\{\alpha \mid (f, \alpha) \in \mathcal{C}_1\} = \alpha_1$ (resp. α_2), then we also have $min\{\omega(\alpha) \mid (f, \omega(\alpha)) \in \mathcal{C}'_1\} = \omega(\alpha_1)$ (resp. $\omega(\alpha_2)$).

Similarly, if $min\{\alpha \mid (f, \alpha) \in \mathcal{C}_1\} > min\{\alpha \mid (f, \alpha) \in \mathcal{C}_2\}$, then $min\{\omega(\alpha) \mid (f, \omega(\alpha)) \in \mathcal{C}'_1\} > min\{\omega(\alpha) \mid (f, \omega(\alpha)) \in \mathcal{C}'_2\}$. Hence, 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 $\omega(.)$, 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 $\omega(.)$, we get $\omega(\alpha) \leq \omega(\beta) = Inc(\mathcal{WA}_2)$. This means $f \notin \pi(\mathcal{WA}_2)$.
Therefore:
$\pi(\mathcal{WA}_1) = \pi(\mathcal{WA}_2)$. ■

Let us illustrate these notions on our running example.

Example 6
Thanks to Lemma 1, in order to compute the repair $\pi(\mathcal{A}_\rhd)$, it is enough to consider only the three bases $\mathcal{WA}_1$, $\mathcal{WA}_2$ and $\mathcal{WA}_3$ as compatible bases of $\mathcal{A}_\rhd$. 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 partially preordered repair is:

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

□

The next section addresses the question of how to compute $\pi(\mathcal{A}_\rhd)$ without enumerating all compatible bases.

Characterization of partially preordered repair

In order to avoid the computation of all assertional bases that are compatible with a partially preordered ABox $\mathcal{A}_\rhd$, we provide a characterization for Definition 5 by introducing the notion of π-accepted assertions. Basically, an assertion is π-accepted if it is strictly preferred to at least one assertion of each assertional conflict of $\mathcal{A}_\rhd$.

Definition 6
Let $\mathbb{L} = (U, \rhd)$ be an uncertainty scale.
Let $\mathcal{K}_\rhd = \langle \mathcal{T}, \mathcal{A}_\rhd, \mathbb{L}\rangle$ be a partially preordered DL-Lite KB.
Let $\mathcal{C}(\mathcal{A}_\rhd)$ denote the set of conflicts of $\mathcal{A}_\rhd$.
An assertion $(f, u) \in \mathcal{A}_\rhd$ is π-accepted iff:

$$\forall \mathcal{C} \in \mathcal{C}(\mathcal{A}_\rhd), \exists (g, v) \in \mathcal{C}, g \neq f, \text{ s.t. } f \rhd g (\text{ i.e., } u \rhd v).$$

Note that the set of assertional conflicts $\mathcal{C}(\mathcal{A}_\rhd)$ is obtained using Definition 1 where the weighted KB $\mathcal{WK}$ and weighted ABox $\mathcal{WA}$ are replaced with the partially preordered KB $\mathcal{K}_\rhd$ and ABox $\mathcal{A}_\rhd$.

Example 7
The set of assertional conflicts of $\mathcal{A}_\rhd$ is:

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

One can easily check 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. Hence they are π-accepted assertions. □

An important result of this paper is that the set of π-accepted assertions corresponds to the repair of the partially preordered ABox $\mathcal{A}_\rhd$ (where weights are omitted).

Proposition 1
An assertion $(f, u) \in \mathcal{A}_\rhd$ is π-accepted iff $f \in \pi(\mathcal{A}_\rhd)$.

Proof:

(i) Assume that $(f, u) \in \mathcal{A}_\rhd$ is π-accepted but $f \notin \pi(\mathcal{A}_\rhd)$. This means that there is a compatible base $\mathcal{WA}$ of $\mathcal{A}_\rhd$ and a weight $\alpha_i \in]0, 1]$ s.t. $(f, \alpha_i) \in \mathcal{WA}$ and $f \notin \pi(\mathcal{WA})$.
Let $Inc(\mathcal{WA}) = \beta$. By Definition 2, this means that $\mathcal{A}^{\geq\beta}$ is inconsistent but $\mathcal{A}^{>\beta}$ is consistent.
Consider a conflict $\{(g, \alpha_j), (h, \alpha_k)\} \in \mathcal{C}(\mathcal{WA})$ where assertions $g, h \in \mathcal{A}^{\geq\beta}$ (such conflict exists since $\mathcal{A}^{\geq\beta}$ is inconsistent). Thus, necessarily $\alpha_j \geq \beta$ and $\alpha_k \geq \beta$ (since $\mathcal{A}^{\geq\beta}$ is inconsistent).
By Definition 3, $f \notin \pi(\mathcal{WA})$ means that $\alpha_i \leq \beta$. Hence $\alpha_j \geq \alpha_i$ and $\alpha_k \geq \alpha_i$. But this contradicts the fact that (f, u) is π-accepted, which ensures that $f \rhd g$ or $f \rhd h$, i.e., $\alpha_i > \alpha_j$ or $\alpha_i > \alpha_k$.

(ii) Assume now that assertion (f, u) is not π-accepted but $f \in \pi(\mathcal{A}_\rhd)$. Assertion (f, u) is not π-accepted means that there is a conflict $\{(g, v), (h, x)\} \in \mathcal{C}(\mathcal{A}_\rhd)$ such that $f \not\rhd g$ and $f \not\rhd h$, i.e., $u \not\rhd v$ and $u \not\rhd x$.
Three distinct cases need to be considered:

  (a) Both $g \rhd f$ and $h \rhd f$ hold, i.e., $v \rhd u$ and $x \rhd u$. This means that in all compatible bases with $\mathcal{A}_\rhd$, both assertions g and h are preferred to f. Let $\mathcal{WA}$ be a compatible base containing (f, α_i), (g, α_j) and (h, α_k), with $\alpha_i, \alpha_j, \alpha_k \in]0, 1]$. Hence $\alpha_j > \alpha_i$ and $\alpha_k > \alpha_i$.
  Assertions g and h are conflicting means that $Inc(\mathcal{WA}) \geq min(\alpha_j, \alpha_k)$. Hence $Inc(\mathcal{WA}) \geq \alpha_i$, thus $f \notin \pi(\mathcal{WA})$. But this contradicts the fact that $f \in \pi(\mathcal{A}_\rhd)$.

(b) Both $f \sim g$ and $h \rhd f$ hold, i.e., $u \sim v$ and $x \rhd u$. In this case, it is enough to have a compatible base $\mathcal{WA}$ containing $(f, \alpha_i), (g, \alpha_j)$ and (h, α_k), with $\alpha_i, \alpha_j, \alpha_k \in]0, 1]$, $\alpha_j > \alpha_i$ and $\alpha_k > \alpha_i$. Such compatible base always exists. Hence, $f \notin \pi(\mathcal{WA})$. But this contradicts the assumption that $f \in \pi(\mathcal{A}_\rhd)$.
Note that the case where $g \rhd f$ but $f \sim h$ is also valid by symmetry.

(c) Both $f \sim g$ and $f \sim h$ hold, i.e., we have $u \sim v$ and $u \sim x$. Then it is enough to have a compatible base $\mathcal{WA}$ containing $(f, \alpha_i), (g, \alpha_j)$ and (h, α_k) where $\alpha_j > \alpha_i$ and $\alpha_k > \alpha_i$. This amounts to case (a) above.

■

Example 8
From Examples 6 and 7, we see that π-accepted assertions (without weights) are exactly those of $\pi(\mathcal{A}_\rhd)$, namely: $\{A(a), A(b), B(c)\}$. □

Properties of partially preordered repair

Thanks to the characterization provided in Definition 6, we are able to state the following results.

Proposition 2

1. *The base $\pi(\mathcal{A}_\rhd)$ is consistent w.r.t. the TBox.*
2. *Computing $\pi(\mathcal{A}_\rhd)$ is done in polynomial time w.r.t. the size of the ABox.*

Proof:

1. The consistency of $\pi(\mathcal{A}_\rhd)$ is straightforward. Since the π-repair $\pi(\mathcal{WA})$ of each compatible base $\mathcal{WA}$ of $\mathcal{A}_\rhd$ is consistent, the intersection of all π-repairs is necessarily consistent.
2. Regarding computational complexity, we recall that computing the set of conflicts $\mathcal{C}(\mathcal{A}_\rhd)$ is done in polynomial time w.r.t. the size of $\mathcal{A}_\rhd$ in DL-Lite. Hence, computing $\pi(\mathcal{A}_\rhd)$ is also done in polynomial time. Indeed, checking if some assertion $(f, u) \in \mathcal{A}_\rhd$ is π-accepted amounts to parsing all assertional conflicts in $\mathcal{C}(\mathcal{A}_\rhd)$. This is done in linear time w.r.t. the size of $\mathcal{C}(\mathcal{A}_\rhd)$ (the size is itself bounded by $\mathcal{O}(|\mathcal{A}_\rhd|^2)$).

■

In addition to the two above results, by construction of $\pi(\mathcal{A}_\rhd)$, it is straightforward to see that when the partial preorder $\rhd$ is a total preorder denoted by $>$, then $\pi(\mathcal{A}_\rhd)$ collapses with the π-repair $\pi(\mathcal{A}_>)$.

We conclude that reasoning (i.e., answering queries) from a partially preordered inconsistent knowledge base amounts to replacing the original ABox $\mathcal{A}_\rhd$ with its repair $\pi(\mathcal{A}_\rhd)$. Indeed, we have established consistency of the repair with respect to the TBox but also tractability of its computation. Furthermore, we have shown that when the preference relation is a total preorder, our method amounts to computing a standard possibilistic repair.

Conclusion

In this paper, we proposed an extension of possibilistic DL-Lite to the case of partially preordered knowledge bases in order to handle inconsistency. The main idea consists in interpreting a partially preordered ABox as a family of compatible weighted ABoxes, then to compute the possibilistic repair of each compatible base and finally to consider the intersection of all possibilistic repairs to produce a single repair for the partially preordered ABox. We proposed a characterization by introducing the notion of π-accepted assertions and showed that the partially preordered repair amounts to computing the set of π-accepted assertions. Most notably, we showed that this computation can be achieved in polynomial time in DL-Lite.

In future work, we plan to investigate methods for enhancing the productivity of the partial repair. For instance, one could consider 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. We expect to show that this is indeed the case by reducing the problem to answering an instance checking query. More generally, we plan to investigate whether methods for computing repairs that are polynomial in the flat and prioritized cases are also polynomial in presence of a partial preorder.

In the context of our research project called AniAge, we plan to apply our findings to the problem of query answering from ontologies representing Southeast Asian dances. The idea is to ask experts in traditional dances to capture the cultural knowledge conveyed by particular dance movements, postures, costumes and props, by semantically enriching dance videos. This is achieved by annotating dance videos w.r.t. the ontology (i.e., the TBox). Experts may assign confidence degrees to their annotations to reflect various reliability levels of the information. This corresponds to defining a priority relation, namely a total preorder, over the assertions of the ABox. However different experts may not share the same meaning of confidence scales. This can be expressed by applying a partial preorder to the ABox. Conflicts may emerge when the same video is annotated differently by several experts. This stresses the importance of handling inconsistency efficiently in order to compute query answers.

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 *12th International Symposium on Frontiers of Combining Systems (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 *Principles of Knowledge Representation and Reasoning (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 *Logic Programming and Nonmonotonic Reasoning (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 *International Joint Conference on Artificial Intelligence (IJCAI), Buenos Aires, Argentina*, 1450–1456.

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 *27th International Joint Conference on Artificial Intelligence, (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.

Dempster, A. P. 1967. Upper and lower probabilities induced by a multivalued mapping. *The Annals of Mathematical Statistics* 38:325–339.

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.

Fagin, R.; Halpern, J. Y.; Moses, Y.; and Vardi, M. Y. 2003. *Reasoning About Knowledge*. MIT Press.

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 *12th International Conference on Principles of Knowledge Representation and Reasoning (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 *12th International Conference on Principles of Knowledge Representation and Reasoning (KR), Toronto, Canada*.

Pan, J. Z.; Stamou, G. B.; Stoilos, G.; and Thomas, E. 2007. Expressive querying over fuzzy DL-Lite ontologies. In *20th 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.

Shafer, G. 1976. *A Mathematical Theory of Evidence*. Princeton University Press.

Spohn, W. 2014. *The Laws of Belief - Ranking Theory and Its Philosophical Applications*. Oxford University Press.

Straccia, U. 2006. Towards top-k query answering in description logics: The case of DL-Lite. In *10th European Conference on Logics in Artificial Intelligence (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. *Journal 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 *Rules and Reasoning - Third International Joint Conference, RuleML+RR, Bolzano, Italy*, 144–158.

Zadeh, L. A. 1978. Fuzzy sets as a basis for a theory of probability. *Fuzzy Sets and Systems* 1:3–28.