Tractable Closure-Based Possibilistic Repair for Partially Ordered DL-Lite Ontologies

Ahmed Laouar[1(✉)], Sihem Belabbes[2], and Salem Benferhat[1]

[1] CRIL, Univ. Artois & CNRS, Lens, France
{laouar,benferhat}@cril.fr
[2] LIASD, IUT de Montreuil, Univ. Paris 8, Saint-Denis, France
belabbes@iut.univ-paris8.fr

Abstract. Inconsistency in formal ontologies is usually addressed by computing repairs for the dataset. There are several strategies for selecting the repairs used to evaluate queries, with various levels of cautiousness and classes of computational complexity. This paper deals with inconsistent partially ordered lightweight ontologies. It introduces a new method that goes beyond the cautious strategies and that is tractable in the possibilistic setting, where uncertainty concerns only the data pieces. The proposed method, called Cπ-repair, proceeds as follows. It first interprets the partially ordered dataset as a family of totally ordered datasets. Then, it computes a single data repair for every totally ordered possibilistic ontology induced from the partially ordered possibilistic ontology. Next, it deductively closes each of these repairs in order to increase their productivity, without introducing conflicts or arbitrary data pieces. Finally, it intersects the closed repairs to obtain a single data repair for the initial ontology. The main contribution of this paper is an equivalent characterization that does not enumerate all the total orders, but also does not suffer from the additional computational cost naturally incurred by the deductive closure. We establish the tractability of our method by reformulating the problem using the notions of dominance and support. Intuitively, the valid conclusions are supported against conflicts by consistent inclusion-minimal subsets of the dataset that dominate all the conflicts. We also study the rationality properties of our method in terms of unconditional and conditional query-answering.

1 Introduction

The Ontology-Based Data Access (OBDA) paradigm relies on an ontology to provide a unified conceptual representation of some domain of interest, in order to improve access to data [27]. Lightweight fragments of description logics such as the DL-Lite family [1,15] are commonly used to encode ontologies, since they allow for efficient query-answering. The conceptual knowledge of an ontology (i.e., the TBox) is usually assumed to be consistent. However, the dataset (i.e., the ABox) may potentially be inconsistent with respect to the TBox. In this case, reasoning with an inconsistent ontology amounts to evaluating queries over one

or several of its data repairs (i.e., inclusion-maximal subsets of the ABox that are consistent with respect to the TBox). A variety of inconsistency-tolerant semantics have been proposed, with different levels of cautiousness and classes of computational complexity. These strategies select one or several repairs of an ABox in order to evaluate queries, with tractability achieved mostly for DL-Lite ontologies (see [3,10] for a survey).

For instance, the well-known Intersection of ABox Repair (IAR) semantics [25] is a cautious strategy. Indeed, it avoids a random selection of the repairs since it evaluates queries over the intersection of all the repairs. It discards all the elements of the ABox that are involved in conflicts (i.e., inclusion-minimal subsets of the ABox that are inconsistent with respect to the TBox). Most notably, it is tractable in DL-Lite.

The Intersection of Closed ABox Repair (ICAR) semantics [25] applies the IAR semantics to the deductive closure of the ABox. This allows to derive more facts from the ABox, so ICAR is more productive[1] than IAR, while it is likewise tractable in DL-Lite. However, ICAR may return undesirable conclusions. Consider a TBox with the concept inclusion axiom Whale $\sqsubseteq$ Mammal, and the concept disjointness axiom Whale $\sqsubseteq \neg$ Shark. Say the ABox contains two facts Whale(Humphrey) and Shark(Humphrey), which are in conflict according to the disjointness axiom. ICAR returns Mammal(Humphrey) as a valid conclusion, although it follows from Whale(Humphrey), which is involved in a conflict.

For prioritized (or totally ordered) ontologies, computing repairs for the ABox is a challenging task. For instance, the Preferred Repair semantics [12] (based on the notion of preferred subtheories [13]) is coNP-complete in data complexity. The grounded repair [11] obtained from a Dung-style argumentation framework is tractable. However, the order relation only applies to the conflicts, so the ABox is assumed to be locally stratified (a similar order defined in propositional logic can be found in [8]). In contrast, three main tractable methods have been identified in [7] for totally ordered ABoxes. The non-defeated repair iteratively applies the IAR semantics to a cumulative sequence of strata of the ABox. The linear-based repair iteratively accumulates a sequence of strata of the ABox, by discarding any stratum that is inconsistent with respect to the TBox or that contradicts the preceding sequence of strata. The possibilistic repair leverages possibility theory, which supports inconsistent reasoning with incomplete, uncertain, qualitative and prioritized information [20,23]. It infers all the facts that are strictly more certain than some inconsistency degree of the ABox, computed from the uncertainty degrees assigned to the elements of the ABox.

A few tractable methods have also been proposed for partially ordered ontologies. There is the Elect method [6] which generalizes both the IAR semantics and the non-defeated repair. There is also the possibilistic repair method defined for partially ordered ontologies [5]. Both Elect and the possibilistic repair method extend the partial order over the ABox into a family of total orders. This yields as many totally ordered ABoxes, for which repairs can be computed then intersected to obtain a single repair for the initial ABox. Such methods are interesting

---

[1] Given two inconsistency-tolerant semantics s_1 and s_2. Then s_1 is more productive than s_2 if any conclusion derived with s_2 can also be derived with s_1.

since they derive the conclusions that follow from all the totally ordered repairs. However, their productivity is hampered by their cautiousness.

A natural way for increasing productivity is to compute the deductive closure of the totally ordered repairs before intersecting them, in the spirit of the Intersection of Closed Repairs (ICR) semantics [9].[2] However, it is well-known that closing the repairs increases time complexity. For instance, for flat ontologies, inference with the IAR semantics and the ICAR semantics is tractable, but it is coNP-complete with the ICR semantics in data complexity [3,9].

In this work, we undertake the difficult task of proposing a more productive, yet efficient method, in the case of partially ordered ontologies. We call the new method Cπ-repair and establish its tractability in the DL-Lite$_{\mathcal{R}}$ fragment, using an interpretation of possibility theory as a strategy for computing the repair. We characterize Cπ-repair equivalently based on the notions of dominance and support. Intuitively, the valid conclusions are supported against conflicts by consistent inclusion-minimal subsets of the ABox that dominate all the conflicts.

We also study the rationality properties of Cπ-repair in terms of unconditional and conditional query-answering mechanisms. In particular, we show that the deductive closure preserves the satisfied properties of possibilistic query-answering from the intersection of the (unclosed) possibilistic repairs.

This paper is organised as follows. Section 2 recalls the basics of DL-Lite$_{\mathcal{R}}$ and the possibilistic repair method. Section 3 defines the closure-based repair method Cπ-repair. Section 4 gives its tractable characterization. Section 5 studies its rationality properties, before concluding.

2   Preliminaries

We recall the underpinnings of the possibilistic repair for partially preordered ontologies [5] that are specified in the DL-Lite$_{\mathcal{R}}$ fragment [15].

Flat (or Non-prioritized) Ontology: A DL-Lite$_{\mathcal{R}}$ ontology is a knowledge base (KB) $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$, where $\mathcal{T}$ is a TBox composed of axioms encoding domain knowledge, and $\mathcal{A}$ is an ABox composed of assertions (i.e., ground facts or data pieces). The axioms may be positive inclusions of concepts (e.g. $B_1 \sqsubseteq B_2$) or of roles (e.g. $R_1 \sqsubseteq R_2$), and they allow to derive new assertions from the ABox. The axioms may also be negative inclusions of concepts (e.g. $B_1 \sqsubseteq \neg B_2$) or of roles (e.g. $R_1 \sqsubseteq \neg R_2$), and they serve to exhibit the conflicts in the ABox.

Definition 1. *(Conflict) Let $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ be a KB. A conflict is a subset $\mathcal{C} \subseteq \mathcal{A}$ such that $\langle \mathcal{T}, \mathcal{C} \rangle$ is inconsistent, and for all $\varphi \in \mathcal{C}$, $\langle \mathcal{T}, \mathcal{C} \setminus \{\varphi\} \rangle$ is consistent. We denote by $\mathsf{Cf}(\mathcal{A})$ the set of all the conflicts of $\mathcal{A}$, a.k.a. the conflict set.*

Here, inconsistency means the absence of a model for the KB (we omit the semantics for space reasons). The conflict set $\mathsf{Cf}(\mathcal{A})$ can be computed in polynomial time in the size of the ABox in DL-Lite$_{\mathcal{R}}$ ontologies [15].

---

[2] Note the difference between the ICR semantics which deductively closes the repairs of an ABox, and the ICAR semantics which computes the repairs for a closed ABox.

Totally Preordered Possibilistic Ontology: Possibilistic logic and possibility the-
ory [4,18,19] are long-standing approaches for reasoning with uncertain informa-
tion, and are closely related to ordinal conditional functions [29] and consonant
belief functions [17,22,28]. Uncertainty can be represented either in extension
using possibility distributions, or in a compact way using weighted logics or
graphical models. Here, we opt for a qualitative representation of the preference
relation induced over the ABox, where only the plausibility ordering between
the assertions matters (see [5] for details and an overview of possibility theory).

Let $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ be a KB, where the axioms in $\mathcal{T}$ are fully certain and free
of conflicts, while the elements of $\mathcal{A}$ may be uncertain and conflicting w.r.t. the
axioms of $\mathcal{T}$. Consider a total preorder $\geq$ over $\mathcal{A}$[3], and let $>$ be the associated
strict order, and let $\equiv$ denote the associated equivalence relation. We denote
the resulting ABox by $\mathcal{A}_\geq$, and it can be represented as a well-ordered partition
$(\mathcal{S}_1, \ldots, \mathcal{S}_n)$ such that:

- $\mathcal{S}_1 \cup \ldots \cup \mathcal{S}_n = \mathcal{A}$.
- $\mathcal{S}_1 = \{\varphi_j \in \mathcal{A} : \text{for all } \varphi_k \in \mathcal{A}, \varphi_j \geq \varphi_k\}$.
- $\mathcal{S}_i = \{\varphi_j \in \mathcal{A} \setminus (\mathcal{S}_1 \cup \cdots \cup \mathcal{S}_{i-1}) : \text{for all } \varphi_k \in \mathcal{A} \setminus (\mathcal{S}_1 \cup \cdots \cup \mathcal{S}_{i-1}), \varphi_j \geq \varphi_k\}$,
  for $i = 2, \ldots, n$.

The assertions in $\mathcal{S}_1$ (resp. $\mathcal{S}_n$) are the most (resp. least) certain. Those in any $\mathcal{S}_i$
$(i = 1, \ldots, n)$ are equally certain.

The totally preordered possibilistic repair [7], denoted $\mathcal{R}(\mathcal{A}_\geq)$, can be com-
puted tractably like so:

- If $\mathcal{K}$ is consistent, then $\mathcal{R}(\mathcal{A}_\geq) = \mathcal{A}$.
- Otherwise, if $\langle \mathcal{T}, \mathcal{S}_1 \rangle$ is inconsistent, then $\mathcal{R}(\mathcal{A}_\geq) = \emptyset$.
- Otherwise, if for some i, $1 \leq i < n$, $\langle \mathcal{T}, \mathcal{S}_1 \cup \ldots \cup \mathcal{S}_i \rangle$ is consistent, and
  $\langle \mathcal{T}, \mathcal{S}_1 \cup \ldots \cup \mathcal{S}_{i+1} \rangle$ is inconsistent, then $\mathcal{R}(\mathcal{A}_\geq) = \mathcal{S}_1 \cup \ldots \cup \mathcal{S}_i$.

Partially Preordered Possibilistic Ontology: Consider a partial preorder $\unrhd$
over $\mathcal{A}$[4]. Let $\rhd$ be the associated strict order, and let $\bowtie$[5] denote incompara-
bility. We denote the resulting ABox by $\mathcal{A}_\unrhd$.

The partially preordered possibilistic repair [5], denoted $\pi(\mathcal{A}_\unrhd)$, relies on the
notion of compatible bases. These are all the totally preordered ABoxes induced
from $\mathcal{A}_\unrhd$ that preserve the ordering between its assertions. A totally preordered
ABox $\mathcal{A}_\geq$ is compatible with $\mathcal{A}_\unrhd$ means that for all $\varphi_j \in \mathcal{A}_\unrhd$, for all $\varphi_k \in \mathcal{A}_\unrhd$,
if $\varphi_j \unrhd \varphi_k$, then $\varphi_j \geq \varphi_k$. This entails that $\varphi_j \bowtie \varphi_k$ extends to three distinct
cases: (i) $\varphi_j > \varphi_k$, (ii) $\varphi_k > \varphi_j$ or (iii) $\varphi_j \equiv \varphi_k$. Thus, $\pi(\mathcal{A}_\unrhd)$ is obtained like so:

1. First, extend the partial preorder $\unrhd$ into a family of total preorders, each of
   which is denoted by $\geq_i$, with $1 \leq i \leq m$ and m is the number of extensions
   of $\unrhd$. Each extension $\geq_i$ defines an ABox $\mathcal{A}_{\geq_i}$ that is compatible with $\mathcal{A}_\unrhd$.

---

[3] A binary relation $\geq$ over $\mathcal{A}$ is a total preorder if it is reflexive and transitive, and
for all $\varphi_j \in \mathcal{A}$, for all $\varphi_k \in \mathcal{A}$, either $\varphi_j \geq \varphi_k$ or $\varphi_k \geq \varphi_j$.

[4] A binary relation $\unrhd$ over $\mathcal{A}$ is a partial preorder if it is reflexive and transitive. Thus
somes elements of $\mathcal{A}$ may be incomparable according to $\unrhd$.

[5] Consider φ_j and φ_k in $\mathcal{A}$. $\varphi_j \bowtie \varphi_k$ means that neither $\varphi_j \unrhd \varphi_k$ nor $\varphi_k \unrhd \varphi_j$ holds.

2. Then for each compatible base $\mathcal{A}_{\geq_i}$, compute its totally preordered possibilistic repair $\mathcal{R}(\mathcal{A}_{\geq_i})$ as defined above.
3. Finally, intersect all the repairs $\mathcal{R}(\mathcal{A}_{\geq_i})$ to obtain $\pi(\mathcal{A}_{\unrhd}) = \bigcap_{i=1}^{m} \mathcal{R}(\mathcal{A}_{\geq_i})$.

The repair $\pi(\mathcal{A}_{\unrhd})$ has been characterized tractably [5], without exhibiting all the extensions $\geq_i$ of $\unrhd$, using the notion of π-accepted assertions.

Definition 2. *(π-accepted assertion) Let $\mathcal{K}_{\unrhd} = \langle \mathcal{T}, \mathcal{A}_{\unrhd} \rangle$ be a partially preordered KB, and $\mathsf{Cf}(\mathcal{A}_{\unrhd})$ its conflict set. An assertion $\varphi_j \in \mathcal{A}_{\unrhd}$ is π-accepted if for all $\mathcal{C} \in \mathsf{Cf}(\mathcal{A}_{\unrhd})$, there is $\varphi_k \in \mathcal{C}, \varphi_j \neq \varphi_k$, s.t. $\varphi_j \rhd \varphi_k$.*

The repair $\pi(\mathcal{A}_{\unrhd})$ is the set of all the π-accepted assertions, and it can be computed in polynomial time in the size of $\mathcal{A}_{\unrhd}$ in DL-Lite$_{\mathcal{R}}$ ontologies [5].

We introduce a toy example of a sales company's information security policy.

Example 1. We build a KB from the following mutually disjoint sets N_C, N_R and N_I, containing respectively concept names, role names and individuals:

- $N_C = \{\mathsf{Reports}, \mathsf{HRfiles}, \mathsf{Manager}, \mathsf{Sales}, \mathsf{Staff}, \mathsf{HR}\}$, where $\mathsf{Reports}$, $\mathsf{HRfiles}$ are file categories and $\mathsf{Manager}$, Sales, Staff and HR are employee positions.
- $N_R = \{\mathsf{Edit}, \mathsf{Sign}, \mathsf{Read}\}$, represent the privileges of an employee on a file.
- $N_I = \{\mathsf{Bob}, \mathsf{Alice}, \mathsf{F17}, \mathsf{F78}\}$, where Bob, Alice are employees and $\mathsf{F17}, \mathsf{F78}$ represent shared files.

Consider a partially preodered KB $\mathcal{K}_{\unrhd} = \langle \mathcal{T}, \mathcal{A}_{\unrhd} \rangle$ depicted in Fig. 1.
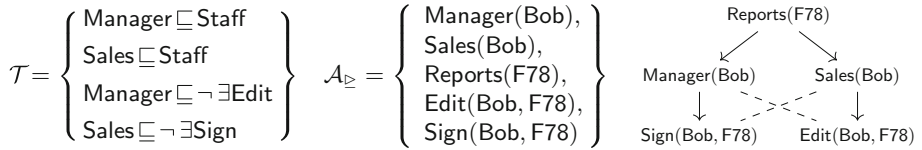

Fig. 1. Left: The TBox $\mathcal{T}$ ($\exists$ indicates the existential restriction on roles). Middle: The ABox $\mathcal{A}_{\unrhd}$. Right: The conflicts of $\mathcal{A}_{\unrhd}$ (dashed lines) and the relation $\unrhd$ (solid arrows represent the strict preference $\rhd$, the other elements are incomparable).

According to $\mathcal{T}$, a manager and a sales person are staff members. A manager (resp. a sales person) does not have editing (resp. signing) rights on files.

Applying Definition 2 to Fig. 1 (right), since $\mathsf{Reports}(\mathsf{F78})$ is strictly preferred to at least one member of each conflict, we get $\pi(\mathcal{A}_{\unrhd}) = \{\mathsf{Reports}(\mathsf{F78})\}$.

□

3   The Cπ-Repair Method

In the literature, the notion of positive deductive closure is a natural way for obtaining more productive repairs. The idea is to apply the positive inclusion axioms of the TBox to the ABox in order to derive new assertions. However, this typically increases the computational cost. In this section, we propose a

new method, called Cπ-repair, which produces a larger partially preordered possibilistic repair, while maintaining tractability in DL-Lite$_\mathcal{R}$ and the satisfiability of the rationality properties. Moreover, the tractability of the proposed method is also applicable in fragments that are more expressive than DL-Lite$_\mathcal{R}$.

Let us first recall the definition of the closure operator in the case of a flat KB [7,15], which can be applied in the partially preordered case (we assume that the individuals included in the closure are limited to those present in the ABox):

Definition 3. *(Closed ABox) Let $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ be a DL-Lite$_\mathcal{R}$ KB. Let $\mathcal{T}_p$ denote the set of all the positive inclusion axioms of $\mathcal{T}$. The deductive closure of $\mathcal{A}$ w.r.t. $\mathcal{T}$ is defined as:*
$cl(\mathcal{A}) = \{B(a)|\langle \mathcal{T}_p, \mathcal{A} \rangle \vDash B(a)$ s.t. B is a concept name in $\mathcal{T}$, a is an individual in $\mathcal{A}\} \bigcup \{R(a,b)|\langle \mathcal{T}_p, \mathcal{A} \rangle \vDash R(a,b)$ s.t. R is a role name in $\mathcal{T}$, a and b are individuals in $\mathcal{A}\}$, where $\vDash$ is the standard DL-Lite$_\mathcal{R}$ inference relation.

For partially preordered ABoxes, the deductive closure can be applied at two different levels to obtain a larger repair. The first option closes the initial ABox $\mathcal{A}_\rhd$ (in the spirit of the ICAR semantics for flat ABoxes [25]). The second option closes each one of the possibilistic repairs $\mathcal{R}(\mathcal{A}_{\geq_i})$ associated with the compatible ABoxes $\mathcal{A}_{\geq_i}$ (in the spirit of the ICR semantics for flat ABoxes [9]).

In the first option, closing the ABox $\mathcal{A}_\rhd$ before computing the repair may lead to undesirable conclusions. For instance, assume a DL-Lite$_\mathcal{R}$ KB where $\mathcal{T} = \{B \sqsubseteq E, A \sqsubseteq \neg B\}$ and $\mathcal{A}_\rhd = \{A(a), B(a)\}$. Then $cl(\mathcal{A}_\rhd) = \{A(a), B(a), E(a)\}$ and $\mathsf{Cf}(cl(\mathcal{A}_\rhd)) = \{\{A(a), B(a)\}\}$. Assume that $E(a) \rhd A(a) \rhd B(a)$. Using Definition 2, both $E(a)$ and $A(a)$ are π-accepted in $cl(\mathcal{A}_\rhd)$. However, including $E(a)$ in the repair is questionable, since it is supported by $B(a)$ which conflicts with the π-accepted assertion $A(a)$. Another issue with this approach concerns the reliability of the assertions that are inferred from incomparable elements. For instance, assume a DL-Lite$_\mathcal{R}$ KB where $\mathcal{T} = \{A \sqsubseteq E, B \sqsubseteq E\}$ and $\mathcal{A}_\rhd = \{A(b), B(b)\}$, where $A(b) \bowtie B(b)$. Then $cl(\mathcal{A}_\rhd) = \{A(b), B(b), E(b)\}$ and all the assertions are π-accepted (since the KB is consistent). The question is which certainty level should be assigned to $E(b)$. The intuition is to consider $E(b)$ to be at least as certain as $A(b)$ and $B(b)$, but this cannot be easily defined in a general way.

In the second option, closing the possibilistic repairs $\mathcal{R}(\mathcal{A}_{\geq_i})$ of the compatible ABoxes $\mathcal{A}_{\geq_i}$ is more appropriate, such that a more productive repair for $\mathcal{A}_\rhd$ can be computed from the intersection of $cl(\mathcal{R}(\mathcal{A}_{\geq_i}))$.

Definition 4. *(Cπ-repair) Let $\mathcal{A}_\rhd$ be a partially preordered ABox. Let $\mathcal{A}_{\geq_i}$, with $1 \leq i \leq m$, denote all its compatible bases and let $\mathcal{R}(\mathcal{A}_{\geq_i})$ be the associated possibilistic repair. The closure-based partially preordered possibilistic repair of $\mathcal{A}_\rhd$, denoted $c\pi(\mathcal{A}_\rhd)$, is obtained as follows:*

$$c\pi(\mathcal{A}_\rhd) = \bigcap_{i=1}^{m} \{cl(\mathcal{R}(\mathcal{A}_{\geq_i})) \mid \mathcal{A}_{\geq_i} \text{ is compatible with } \mathcal{A}_\rhd\}.$$

The closure-based repair $c\pi(\mathcal{A}_\rhd)$ computed with this method is more productive than both the repair $\pi(\mathcal{A}_\rhd)$ and its closure $cl(\pi(\mathcal{A}_\rhd))$. Namely:

$$\pi(\mathcal{A}_\rhd) \subseteq cl(\pi(\mathcal{A}_\rhd)) \subseteq c\pi(\mathcal{A}_\rhd).$$

Figure 2 shows the Cπ-repair process of Definition 4.

- First, extend the partially preordered ABox $\mathcal{A}_{\unrhd}$ into a family of totally preordered ABoxes, denoted $A_{\geq_1}, \ldots, A_{\geq_m}$ ("Extend" arc).
- Then, compute the repair $\mathcal{R}(A_{\geq_i})$ of each compatible base ("Repair" arc), and its deductive closure ("Close" arc).
- Finally, intersect all the closed repairs $cl(\mathcal{R}(A_{\geq_i}))$ ("Intersect" arc) to obtain a single closure-based repair for the initial ABox $c\pi(\mathcal{A}_{\unrhd})$.
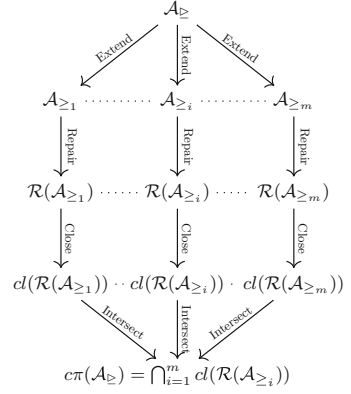

Fig. 2. The Cπ-repair process

Obviously, this method is naive and computationally expensive since enumerating all the compatible ABoxes may be exponential in the worst case.

Example 2. From Example 1, we have: $\pi(\mathcal{A}_{\unrhd}) = cl(\pi(\mathcal{A}_{\unrhd})) = \{\mathsf{Reports(F78)}\}$. Figure 1 illustrates $\mathcal{A}_{\unrhd}$, its two conflicts and the relation $\unrhd$. Recall that any compatible base $\mathcal{A}_{\geq_i}$ preserves the ordering between the assertions of $\mathcal{A}_{\unrhd}$. Since $\mathsf{Reports(F78)}$ is the most certain assertion in $\mathcal{A}_{\unrhd}$, it belongs to every $\mathcal{R}(\mathcal{A}_{\geq_i})$. Moreover, it is easy to check that each $\mathcal{R}(\mathcal{A}_{\geq_i})$ contains either $\mathsf{Manager(Bob)}$ or $\mathsf{Sales(Bob)}$. Using the axioms $\mathsf{Manager} \sqsubseteq \mathsf{Staff}$ and $\mathsf{Sales} \sqsubseteq \mathsf{Staff}$ of $\mathcal{T}$, one can infer $\mathsf{Staff(Bob)}$ from each closed repair $cl(\mathcal{R}(\mathcal{A}_{\geq_i}))$. Therefore, $c\pi(\mathcal{A}_{\unrhd}) = \{\mathsf{Reports(F78)}, \mathsf{Staff(Bob)}\}$. So, $c\pi(\mathcal{A}_{\unrhd})$ is larger than both $\pi(\mathcal{A}_{\unrhd})$ and $cl(\pi(\mathcal{A}_{\unrhd}))$. □

4   Characterization of Cπ-Repair

In this section, we propose an equivalent characterization of the closure-based repair introduced in Definition 4 without enumerating all the compatible bases. We first introduce two notions called support and dominance. The support (or an argument) of an assertion is an inclusion-minimal consistent subset of the ABox that allows to derive it. Formally:

Definition 5. *(Support) Let $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ be a DL-Lite$_{\mathcal{R}}$ KB. Let B be a concept name and R be a role name in $\mathcal{T}$. Let a and b be individuals in $\mathcal{A}$. The subset $\mathcal{B} \subseteq \mathcal{A}$ is a support for $B(a)$ (resp. $R(a, b)$) in $\mathcal{A}$ if:*

- $\langle \mathcal{T}, \mathcal{B} \rangle$ *is consistent, and*
- $\langle \mathcal{T}, \mathcal{B} \rangle \vDash B(a)$, *(resp.* $\langle \mathcal{T}, \mathcal{B} \rangle \vDash R(a, b)$*), and*
- *for all* $\mathcal{B}' \subsetneq \mathcal{B}$, $\langle \mathcal{T}, \mathcal{B}' \rangle \nvDash B(a)$, *(resp.* $\langle \mathcal{T}, \mathcal{B}' \rangle \nvDash R(a, b)$*).*

Where $\vDash$ is the standard DL-Lite$_{\mathcal{R}}$ inference relation.

Example 3. We continue Example 1. We have $\{\mathsf{Manager(Bob)}\}$ supports $\mathsf{Staff(Bob)}$ since $\langle \mathcal{T}, \{\mathsf{Manager(Bob)}\} \rangle \vDash \mathsf{Staff(Bob)}$, and it is minimal and consistent. □

The dominance relation is a way for extending the partial preorder defined over an ABox into a partial preorder defined over the subsets of the ABox. Such extension allows to compare supports and conflicts. Intuitively, the dominance of a partially preordered subset over another requires that each element of the former be strictly more certain than at least one element of the latter. Formally:

Definition 6. *(Dominance) Let $\mathcal{K}_{\unrhd} = \langle \mathcal{T}, \mathcal{A}_{\unrhd} \rangle$ be a partially preordered KB equipped with $\unrhd$. Let $\mathcal{B}_1 \subseteq \mathcal{A}_{\unrhd}$ and $\mathcal{B}_2 \subseteq \mathcal{A}_{\unrhd}$. We say that $\mathcal{B}_1$ dominates $\mathcal{B}_2$, denoted $\mathcal{B}_1 \rhd^{dom} \mathcal{B}_2$, if: for all $\varphi_j \in \mathcal{B}_1$, there is $\varphi_k \in \mathcal{B}_2$ s.t. $\varphi_j \rhd \varphi_k$.*

Example 4. Let $\mathcal{B}_1$, $\mathcal{B}_2$ and $\mathcal{B}_3$ be three subsets of $\mathcal{A}_{\unrhd}$ of Example 1, as illustrated by Fig. 3. $\mathcal{B}_1 \rhd^{dom} \mathcal{B}_2$ holds because Reports(F78) $\rhd$ Sales(Bob) and Manager(Bob) $\rhd$ Sign(Bob, F78). $\mathcal{B}_1 \rhd^{dom} \mathcal{B}_3$ does not hold because Manager(Bob) $\not\rhd$ Sales(Bob) and Manager(Bob) $\not\rhd$ Edit(Bob, F78).

$\mathcal{B}_1$: Reports(F78) $\longrightarrow$ Manager(Bob)   Reports(F78) $\longrightarrow$ Manager(Bob) : $\mathcal{B}_1$

$\mathcal{B}_2$: Sales(Bob)   Sign(Bob, F78)   Sales(Bob) $\longrightarrow$ Edit(Bob, F78) : $\mathcal{B}_3$

(a) $\mathcal{B}_1 \rhd^{dom} \mathcal{B}_2$ holds.   (b) $\mathcal{B}_1 \rhd^{dom} \mathcal{B}_3$ does not hold.
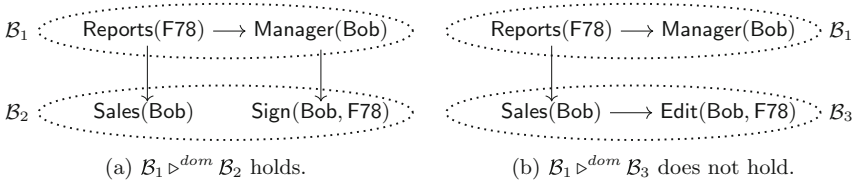
Fig. 3. Solid arrows represent the strict preference.

□

Before characterizing Cπ-repair in general, we first discuss the special case where the ABox is consistent w.r.t. the TBox, i.e., the conflict set is empty. Hence, Cπ-repair simply amounts to applying standard DL-Lite$_{\mathcal{R}}$ inference. Formally:

Lemma 1. *Let $\mathcal{K}_{\unrhd} = \langle \mathcal{T}, \mathcal{A}_{\unrhd} \rangle$ be a consistent, partially preordered KB, i.e., $\mathsf{Cf}(\mathcal{A}_{\unrhd}) = \emptyset$. Consider $cl(\cdot)$ given by Definition 3. Let φ be an assertion. Then: $c\pi(\mathcal{A}_{\unrhd}) = cl(\mathcal{A}_{\unrhd})$. Equivalently: $\varphi \in c\pi(\mathcal{A}_{\unrhd})$ iff there is $\mathcal{B} \subseteq \mathcal{A}_{\unrhd}$ s.t. $\langle \mathcal{T}, \mathcal{B} \rangle \vDash \varphi$.*

In the rest of this paper, we focus on the case where the ABox is inconsistent w.r.t. the TBox. Our goal is to use the notions of dominance and support (see Definitions 5 and 6) to characterize equivalently the assertions in the Cπ-repair. This allows to avoid enumerating all the totally preordered extensions of $\mathcal{A}_{\unrhd}$. The idea is that an assertion φ belongs to Cπ-repair if and only if for every conflict $\mathcal{C}$ in the ABox, there is a support $\mathcal{B}$ of φ that dominates $\mathcal{C}$. We provide two propositions to confirm this intuitive characterization.

Proposition 1. *Let $\mathcal{K}_{\unrhd} = \langle \mathcal{T}, \mathcal{A}_{\unrhd} \rangle$ be an inconsistent, partially preordered KB. Let $\mathsf{Cf}(\mathcal{A}_{\unrhd})$ be its conflict set and let φ be an assertion. If for all $\mathcal{C} \in \mathsf{Cf}(\mathcal{A}_{\unrhd})$, there is $\mathcal{B} \subseteq \mathcal{A}_{\unrhd}$ s.t.:*

1. $\mathcal{B}$ supports φ (as per Definition 5), and
2. $\mathcal{B} \rhd^{dom} \mathcal{C}$ (as per Definition 6),

then $\varphi \in c\pi(\mathcal{A}_{\unrhd})$.

We illustrate this result with our running example.

Example 5. Consider again Example 2. Let us use Proposition 1 to check that the assertion $\mathsf{Staff(Bob)}$ is indeed in $c\pi(\mathcal{A}_{\unrhd})$. For each conflict in $\mathsf{Cf}(\mathcal{A}_{\unrhd}) = \{\{\mathsf{Manager(Bob)}, \mathsf{Edit(Bob, F78)}\}, \{\mathsf{Sales(Bob)}, \mathsf{Sign(Bob, F78)}\}\}$, it suffices to exhibit a dominating support for $\mathsf{Staff(Bob)}$, like so:

- For the conflict $\mathcal{C}_1 = \{\mathsf{Manager(Bob)}, \mathsf{Edit(Bob, F78)}\}$, $\mathcal{B}_1 = \{\mathsf{Sales(Bob)}\}$ supports $\mathsf{Staff(Bob)}$ and $\mathcal{B}_1 \rhd^{dom} \mathcal{C}_1$ (since $\mathsf{Sales(Bob)} \rhd \mathsf{Edit(Bob, F78)}$).
- For the conflict $\mathcal{C}_2 = \{\mathsf{Sales(Bob)}, \mathsf{Sign(Bob, F78)}\}$, $\mathcal{B}_2 = \{\mathsf{Manager(Bob)}\}$ supports $\mathsf{Staff(Bob)}$ and $\mathcal{B}_2 \rhd^{dom} \mathcal{C}_2$ (since $\mathsf{Manager(Bob)} \rhd \mathsf{Sign(Bob, F78)}$).

$\square$

The other direction of Proposition 1, given in Proposition 2, is also true. In particular, if the characterization "for every conflict $\mathcal{C}$ in $\mathsf{Cf}(\mathcal{A}_{\unrhd})$, there is a support $\mathcal{B}$ of φ in $\mathcal{A}_{\unrhd}$ that dominates $\mathcal{C}$" is not true, then φ cannot belong to $\mathrm{C}\pi$-repair. The proposition also covers the particular case of an assertion without a support, which cannot belong to $\mathrm{C}\pi$-repair.

Proposition 2. *Let $\mathcal{K}_{\unrhd} = \langle \mathcal{T}, \mathcal{A}_{\unrhd} \rangle$ be an inconsistent, partially preordered KB. Let $\mathsf{Cf}(\mathcal{A}_{\unrhd})$ be its conflict set and φ be an assertion.*
If $\varphi \in c\pi(\mathcal{A}_{\unrhd})$, then for all $\mathcal{C} \in \mathsf{Cf}(\mathcal{A}_{\unrhd})$, there is $\mathcal{B} \subseteq \mathcal{A}_{\unrhd}$ s.t.:

1. $\mathcal{B}$ supports φ (as per Definition 5), and
2. $\mathcal{B} \rhd^{dom} \mathcal{C}$ (as per Definition 6).

We illustrate this result with our running example.

Example 6. Let $\mathcal{K}'_{\unrhd} = \langle \mathcal{T}', \mathcal{A}_{\unrhd} \rangle$, where the TBox $\mathcal{T}'$ is:

$$\mathcal{T}' = \{\mathsf{Manager} \sqsubseteq \neg \exists \mathsf{Edit}, \mathsf{Sales} \sqsubseteq \neg \exists \mathsf{Sign}, \mathsf{Sales} \sqsubseteq \mathsf{Staff}, \exists \mathsf{Edit} \sqsubseteq \mathsf{Staff}\}.$$

Thus, a manager (resp. a sales person) does not have editing (resp. signing) rights, and a sales person and a person with editing rights are staff members. The ABox $\mathcal{A}_{\unrhd}$ is the one of Example 1.

Consider $\mathcal{A}_{\geq_1}$ and $\mathcal{A}_{\geq_2}$ two ABoxes compatible with $\mathcal{A}_{\unrhd}$, and their well-ordered partitions, $\mathcal{A}_{\geq_1} = (\mathcal{S}_1 \cup \mathcal{S}_2 \cup \mathcal{S}_3 \cup \mathcal{S}_4)$ and $\mathcal{A}_{\geq_2} = (\mathcal{S}_1 \cup \mathcal{S}_2 \cup \mathcal{S}'_3 \cup \mathcal{S}'_4)$ such that: $\mathcal{S}_1 = \{\mathsf{Reports(F78)}\}$, $\mathcal{S}_2 = \{\mathsf{Manager(Bob)}\}$, $\mathcal{S}_3 = \{\mathsf{Sales(Bob)}\}$, $\mathcal{S}'_3 = \{\mathsf{Sales(Bob)}, \mathsf{Sign(Bob, F78)}\}$, $\mathcal{S}_4 = \{\mathsf{Sign(Bob, F78)}, \mathsf{Edit(Bob, F78)}\}$, and $\mathcal{S}'_4 = \{\mathsf{Edit(Bob, F78)}\}$.

Figure 4 illustrates $\mathcal{A}_{\unrhd}$, $\mathcal{A}_{\geq_1}$ and $\mathcal{A}_{\geq_2}$ (recall that $\equiv$ denotes equal certainty). It is easy to check that $\mathcal{A}_{\geq_1}$ and $\mathcal{A}_{\geq_2}$ are compatible with $\mathcal{A}_{\unrhd}$. Their associated repairs are: $\mathcal{R}(\mathcal{A}_{\geq_1}) = \{\mathsf{Reports(F78)}, \mathsf{Manager(Bob)}, \mathsf{Sales(Bob)}\}$ and $\mathcal{R}(\mathcal{A}_{\geq_2}) = \{\mathsf{Reports(F78)}, \mathsf{Manager(Bob)}\}$.

Consider the assertion $\mathsf{Staff(Bob)}$ and its two supports, $\mathcal{B}_1 = \{\mathsf{Sales(Bob)}\}$ and $\mathcal{B}_2 = \{\mathsf{Edit(Bob, F78)}\}$. Notice that $\mathcal{R}(\mathcal{A}_{\geq_1}) \models \mathsf{Staff(Bob)}$ but $\mathcal{R}(\mathcal{A}_{\geq_2}) \not\models \mathsf{Staff(Bob)}$. Hence, $\mathsf{Staff(Bob)} \notin c\pi(\mathcal{A}_{\unrhd})$. Proposition 2 confirms this result, since neither $\mathcal{B}_1$ nor $\mathcal{B}_2$ dominates the conflict $\{\mathsf{Sales(Bob)}, \mathsf{Sign(Bob, F78)}\}$.
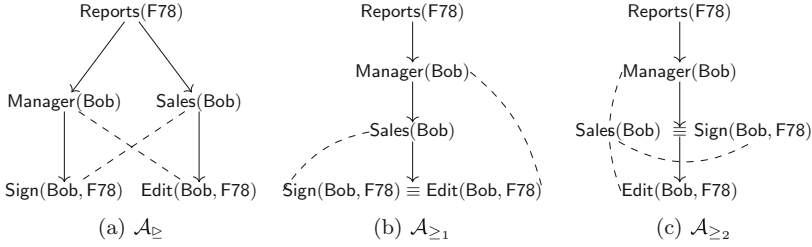
$\square$

Fig. 4. Solid arrows depict strict preference. Dashed lines show the conflicts.

Propositions 1 and 2 provide a full characterization for membership in $c\pi(\mathcal{A}_{\unrhd})$ based on the notions of support and dominance. The next proposition states (as expected) that $c\pi(\mathcal{A}_{\unrhd})$ is consistent and more productive than $\pi(\mathcal{A}_{\unrhd})$. Formally:

Proposition 3.

1. $\langle \mathcal{T}, c\pi(\mathcal{A}_{\unrhd}) \rangle$ is consistent.
2. $\pi(\mathcal{A}_{\unrhd}) \subseteq c\pi(\mathcal{A}_{\unrhd})$. The converse is false (i.e., $c\pi(\mathcal{A}_{\unrhd}) \not\subset \pi(\mathcal{A}_{\unrhd})$).

Example 2 confirms that $c\pi(\mathcal{A}_{\unrhd}) \not\subset \pi(\mathcal{A}_{\unrhd})$.

The next proposition establishes the tractability of $c\pi(\mathcal{A}_{\unrhd})$. This follows from the characterization given in Propositions 1 and 2, i.e., using the notions of dominance and support. Indeed, computing the conflicts and the supports can be achieved in polynomial time [10]. Besides, it can be shown that the number of conflicts and supports is bounded by $|cln(\mathcal{T})| * |\mathcal{A}_{\unrhd}|$ (where $cln(\mathcal{T})$ denotes the negative closure of the TBox $\mathcal{T}$, i.e., all the negative axioms that can be inferred from it). Moreover, in the context of OBDA, the size of the TBox is often considered negligible compared to the size of the ABox, thus the main focus is on data complexity. Lastly, it is important to note that retrieving all the conflicts beforehand is not required. Instead, checking whether an assertion is in Cπ-repair can be performed by progressively examining the conflicts (an implementation is available at https://github.com/ahmedlaouar/py_reasoner). This incremental feature is particularly beneficial for evolving ABoxes.

Proposition 4. Let $\mathcal{K}_{\unrhd} = \langle \mathcal{T}, \mathcal{A}_{\unrhd} \rangle$ be a partially preordered KB and φ be an assertion. Checking if $\varphi \in c\pi(\mathcal{A}_{\unrhd})$ is done in polynomial time in DL-Lite$_{\mathcal{R}}$.

5 Rationality Properties of π-Acceptance and Cπ-Repair

In this section, we study the rationality properties of query-answering using the possibilistic repair method and its closure-based version.

Let $\mathcal{K}_{\unrhd} = \langle \mathcal{T}, \mathcal{A}_{\unrhd} \rangle$ be a partially preordered KB which may be inconsistent and let q be a query. Consider the KB's possibilistic repairs $\pi(\mathcal{A}_{\unrhd})$ (Definition 2) and $c\pi(\mathcal{A}_{\unrhd})$ (Definition 4).

Let us start with unconditional query-answering, which amounts to checking whether the query q follows from the repair $\pi(\mathcal{A}_{\unrhd})$ (resp. $c\pi(\mathcal{A}_{\unrhd})$), denoted with the symbol $\vDash^{\pi}$ (resp. $\vDash^{c\pi}$), and $\vDash$ denotes standard DL-Lite$_{\mathcal{R}}$ inference. Formally:

$$\mathcal{K}_{\unrhd} \vDash^{\pi} q \text{ (resp. } \mathcal{K}_{\unrhd} \vDash^{c\pi} q) \text{ iff } \langle \mathcal{T}, \pi(\mathcal{A}_{\unrhd}) \rangle \vDash q \text{ (resp. } \langle \mathcal{T}, c\pi(\mathcal{A}_{\unrhd}) \rangle \vDash q) \qquad (1)$$

The following result states that the unconditional inferences $\vDash^{\pi}$ and $\vDash^{c\pi}$ meet the rationality properties of unconditional inconsistency-tolerant semantics defined in [2]. Namely:

Proposition 5. *The unconditional possibilistic inference relation $\vDash^s$ (with $s \in \{\pi, c\pi\}$) satisfies the following properties:*

- **QCE** *(Query Conjunction Elimination) If $\mathcal{K}_{\unrhd} \vDash^s q_1 \wedge q_2$ then $\mathcal{K}_{\unrhd} \vDash^s q_1$ and $\mathcal{K}_{\unrhd} \vDash^s q_2$.*
- **QCI** *(Query Conjunction Introduction) If $\mathcal{K}_{\unrhd} \vDash^s q_1$ and $\mathcal{K}_{\unrhd} \vDash^s q_2$ then $\mathcal{K}_{\unrhd} \vDash^s q_1 \wedge q_2$.*
- **Cons** *(Consistency) For any set of assertions $\mathcal{B}$, if $\mathcal{K}_{\unrhd} \vDash^s \mathcal{B}$ then $\langle \mathcal{T}, \mathcal{B} \rangle$ is consistent.*
- **ConsC** *(Consistency of Conjunction) For any set of assertions $\mathcal{B}$, if for all $\varphi \in \mathcal{B}$, $\mathcal{K}_{\unrhd} \vDash^s \varphi$, then $\langle \mathcal{T}, \mathcal{B} \rangle$ is consistent.*
- **ConsS** *(Consistency of Support) For any set of assertions $\mathcal{B}$, if $\mathcal{K}_{\unrhd} \vDash^s \mathcal{B}$ then there is a maximally consistent subset $\mathcal{A}'$ of $\mathcal{A}_{\unrhd}$ s.t. $\langle \mathcal{T}, \mathcal{A}' \rangle \vDash \mathcal{B}$.*

The proof of Proposition 5 is immediate since it is based on a direct application of standard DL-Lite entailment to the repairs $\pi(\mathcal{A}_{\unrhd})$ (resp. $c\pi(\mathcal{A}_{\unrhd})$).

We now focus on conditional query-answering, which amounts to querying a partially preordered KB under a given set of assertions considered fully reliable and consistent with respect to the TBox, called an observation or a fully observable set and denoted by $\mathcal{O}$. We write $\mathcal{O} \mathrel{\vdash}^s_{\mathcal{K}_{\unrhd}} q$ to indicate that q follows from the KB $\mathcal{K}_{\unrhd}$, under the observation $\mathcal{O}$, using the inconsistency-tolerant semantics s (here $s = \pi$ for the possibilistic repair and $s = c\pi$ for its closure-based version). A standard way to proceed is to first add $\mathcal{O}$ to the ABox with the highest priority, then apply the possibilistic repair method (and its closure-based version) using Eq. 1 to unconditionally answer queries from the augmented KB.

Let us denote by $\mathcal{K}_{\mathcal{O}} = \langle \mathcal{T}, \mathcal{A}_{\unrhd_{\mathcal{O}}} \rangle$ the augmented KB where $\mathcal{A}_{\unrhd_{\mathcal{O}}} = (\mathcal{A}_{\unrhd} \cup \mathcal{O})$ results from adding $\mathcal{O}$ to $\mathcal{A}_{\unrhd}$ with the highest priority. Moreover, the partial preorder $\unrhd_{\mathcal{O}}$ over $(\mathcal{A}_{\unrhd} \cup \mathcal{O})$ is obtained from $\unrhd$ as follows:

(i) For all $\varphi_1 \in \mathcal{O}$, for all $\varphi_2 \in \mathcal{O}$: $\varphi_1 \unrhd_{\mathcal{O}} \varphi_2$ and $\varphi_2 \unrhd_{\mathcal{O}} \varphi_1$ (i.e., φ_1 and φ_2 are equally reliable).
(ii) For all $\varphi_1 \in \mathcal{O}$, for all $\varphi_2 \in \mathcal{A}_{\unrhd} \setminus \mathcal{O}$: $\varphi_1 \rhd_{\mathcal{O}} \varphi_2$ (i.e., every $\varphi_1 \in \mathcal{O}$ is strictly more preferred than any $\varphi_2 \in \mathcal{A}_{\unrhd} \setminus \mathcal{O}$. This serves to give priority to $\mathcal{O}$).
(iii) For all $\varphi_1 \in \mathcal{A}_{\unrhd} \setminus \mathcal{O}$, for all $\varphi_2 \in \mathcal{A}_{\unrhd} \setminus \mathcal{O}$: $\varphi_1 \unrhd_{\mathcal{O}} \varphi_2$ iff $\varphi_1 \unrhd \varphi_2$ (i.e., the relative ordering between the elements of $\mathcal{A}_{\unrhd}$ that are not in $\mathcal{O}$ is preserved).

Next, we define the partially preordered conditional query-answering relation.

Definition 7. *(Conditional inference) Let $\mathcal{K}_{\unrhd} = \langle \mathcal{T}, \mathcal{A}_{\unrhd} \rangle$ be a partially pre-ordered KB, $\mathcal{O}$ an observation and q a query. Then q follows from $\mathcal{K}_{\unrhd}$ and $\mathcal{O}$, denoted $\mathcal{O} \vdash^s_{\mathcal{K}_{\unrhd}} q$, if $\mathcal{K}_{\mathcal{O}} \vDash^s q$ (with $s \in \{\pi, c\pi\}$), where $\mathcal{K}_{\mathcal{O}} = \langle \mathcal{T}, \mathcal{A}_{\unrhd_{\mathcal{O}}} \rangle$ is the augmented KB and $\unrhd_{\mathcal{O}}$ its associated partial preorder (described in (i), (ii), (iii)), and $\vDash^s$ is the unconditional query-answering relation given by Eq. 1.*

One can check that $\mathcal{O} \vdash^s_{\mathcal{K}_{\unrhd}} q$ is non-monotonic for both semantics (the possibilistic repair and its closure-based version). The well-known System P [24], originally defined in the context of propositional logic, has been adapted to DL-Lite$_{\mathcal{R}}$ in [3] (see also [14,21] for an adaptation to richer description logics).

The adaptation of System P's rules is given below, where $\mathcal{K}_{\unrhd} = \langle \mathcal{T}, \mathcal{A}_{\unrhd} \rangle$ is a KB, $\mathcal{O}_1, \mathcal{O}_2, \mathcal{O}_3$ are observations, s is an inconsistency-tolerant semantics with $s \in \{\pi, c\pi\}$, $\vDash$ and $\equiv$ denote standard DL-Lite$_{\mathcal{R}}$ inference and equivalence:

- **R** (Reflexivity) $\mathcal{O}_1 \vdash^s_{\mathcal{K}_{\unrhd}} \mathcal{O}_1$.
- **LLE** (Left Logical Equivalence) If $\langle \mathcal{T}, \mathcal{O}_1 \rangle \equiv \langle \mathcal{T}, \mathcal{O}_2 \rangle$ and $\mathcal{O}_1 \vdash^s_{\mathcal{K}_{\unrhd}} \mathcal{O}_3$ then $\mathcal{O}_2 \vdash^s_{\mathcal{K}_{\unrhd}} \mathcal{O}_3$.
- **RW** (Right Weakening) If $\langle \mathcal{T}, \mathcal{O}_1 \rangle \vDash \langle \mathcal{T}, \mathcal{O}_2 \rangle$ and $\mathcal{O}_3 \vdash^s_{\mathcal{K}_{\unrhd}} \mathcal{O}_1$, then $\mathcal{O}_3 \vdash^s_{\mathcal{K}_{\unrhd}} \mathcal{O}_2$.
- **Cut** If $\mathcal{O}_1 \vdash^s_{\mathcal{K}_{\unrhd}} \mathcal{O}_2$ and $\mathcal{O}_1 \cup \mathcal{O}_2 \vdash^s_{\mathcal{K}_{\unrhd}} \mathcal{O}_3$, then $\mathcal{O}_1 \vdash^s_{\mathcal{K}_{\unrhd}} \mathcal{O}_3$.
- **CM** (Cautious Monotony) If $\mathcal{O}_1 \vdash^s_{\mathcal{K}_{\unrhd}} \mathcal{O}_2$ and $\mathcal{O}_1 \vdash^s_{\mathcal{K}_{\unrhd}} \mathcal{O}_3$, then $\mathcal{O}_1 \cup \mathcal{O}_2 \vdash^s_{\mathcal{K}_{\unrhd}} \mathcal{O}_3$.
- **And** If $\mathcal{O}_1 \vdash^s_{\mathcal{K}_{\unrhd}} \mathcal{O}_2$ and $\mathcal{O}_1 \vdash^s_{\mathcal{K}_{\unrhd}} \mathcal{O}_3$, then $\mathcal{O}_1 \vdash^s_{\mathcal{K}_{\unrhd}} \mathcal{O}_2 \cup \mathcal{O}_3$.

In this paper, we propose to also consider two additional properties, originally defined in propositional logic, and which go beyond cautious monotony:

- **RM** (Rational Monotony) If $\mathcal{O}_1 \vdash^s_{\mathcal{K}_{\unrhd}} \mathcal{O}_3$, then $\mathcal{O}_1 \cup \mathcal{O}_2 \vdash^s_{\mathcal{K}_{\unrhd}} \mathcal{O}_3$ or $\langle \mathcal{T}, (\mathcal{O}_1 \cup \mathcal{O}_2 \cup \mathcal{A}_{\unrhd}) \rangle$ is inconsistent.
- **Comp** (Completeness) If $\mathcal{O}_1 \vdash^s_{\mathcal{K}_{\unrhd}} \mathcal{O}_3$, then either $\mathcal{O}_1 \cup \mathcal{O}_2 \vdash^s_{\mathcal{K}_{\unrhd}} \mathcal{O}_3$ or $\langle \mathcal{T}, (\mathcal{O}_1 \cup \mathcal{O}_2 \cup \mathcal{O}_3 \cup \mathcal{A}_{\unrhd}) \rangle$ is inconsistent.

Note that the adaptation of the last two properties that we propose uses the notion of inconsistency instead of negation in the original definition of rational monotony, and uses a disjunctive interpretation of **RM**[6]. Here, **RM** states that given a new observation, we can continue to believe in the previous plausible consequences of the KB, or the new observation conflicts with the KB. The **Comp** rule[7] is stronger than **RM**, and states that given a new observation $\mathcal{O}_2$, then either $\mathcal{O}_3$ continues to be derived from both $\mathcal{O}_1$ and $\mathcal{O}_2$, or $\mathcal{O}_2$ contradicts the whole KB (plus itself). Intuitively, this means that either we continue to believe in $\mathcal{O}_3$, or we should believe in its negation (there is no room for ignoring $\mathcal{O}_3$).

The next proposition summarizes the results of the conditional properties:

---

[6] In propositional logic (PL), **RM** is defined as: if $\alpha \vdash_{\mathcal{K}} \gamma$ and $\alpha \not\vdash_{\mathcal{K}} \neg\beta$, then $\alpha \wedge \beta \vdash_{\mathcal{K}} \gamma$, where α, β and γ are PL formulas. Our adaptation consists first in rewriting **RM** equivalently in a disjunctive way: if $\alpha \vdash_{\mathcal{K}} \gamma$, then $\alpha \wedge \beta \vdash_{\mathcal{K}} \gamma$ or $\alpha \vdash_{\mathcal{K}} \neg\beta$. Lastly, we replace $\alpha \vdash_{\mathcal{K}} \neg\beta$ with $\alpha \wedge \beta$ is inconsistent with the KB.

[7] In PL, **Comp** is defined as: if $\alpha \vdash_{\mathcal{K}} \gamma$, then either $\alpha \wedge \beta \vdash_{\mathcal{K}} \gamma$ or $\alpha \wedge \beta \vdash_{\mathcal{K}} \neg\gamma$. Here, we simply replace $\alpha \wedge \beta \vdash_{\mathcal{K}} \neg\gamma$ with $\alpha \wedge \beta \wedge \gamma$ is inconsistent with the KB.

Proposition 6. *Let $\mathcal{K}_{\trianglerighteq} = \langle \mathcal{T}, \mathcal{A}_{\trianglerighteq} \rangle$ be a partially preordered KB, $\mathcal{O}$ be an observation and q be a query. The query-answering relations $\vdash^{\pi}_{\mathcal{K}_{\trianglerighteq}}$ and $\vdash^{c\pi}_{\mathcal{K}_{\trianglerighteq}}$ satisfy the properties* **R**, **LLE**, **RW**, **Cut**, **CM** *and* **And**. *However, they fail to satisfy* **RM** *and* **Comp**.

6   Concluding Discussions

Developing tractable and safe methods for inconsistency management is a challenge and is crucial for dealing with inconsistent large-scale knowledge bases. This paper follows this research line where we tackled the issue of computing a productive repair for possibilistic partially preordered ontologies. We defined the Cπ-repair method which interprets a partial preorder into a family of compatible ABoxes, computes their possibilistic repairs, closes those repairs and intersects them to yield a more productive repair.

An important result of this paper is that we characterized this method equivalently using the notions of dominance and support, which ensures the tractable calculation of the repair. This characterization can be generalized easily to more expressive description languages (it suffices to replace the DL-Lite$_{\mathcal{R}}$ inference relation in the support definition by that of a more expressive language). However, tractability is guaranteed only if the computation of the conflicts and supports is performed in polynomial time and their size remains polynomial in the ABoxe's size (a detailed discussion is given below). A future work is to characterize the linear repair [26] and the Elect [6] methods to partially preordered ABoxes.

We conclude this paper with a few discussion points on the rational properties as well as on the possibility of generalizing our method to richer languages or other inconsistency-tolerant semantics.

On the Rational Properties: The two possibilistic semantics studied in this paper satisfy the unconditional properties (Proposition 5) and the rules of System P (Proposition 6). If these propositions seem natural, even minimal, they are not always satisfied by some inconsistency-tolerant semantics. For example, the so-called majority semantics (a query is valid if it is obtained from the majority of the repairs of an inconsistent ABox) does not satisfy these minimal properties. More precisely, in [3] it has been shown that majority-based inference does not satisfy Cut, Cautious Monotony, and And properties, even in DL-Lite$_{\mathcal{R}}$. Another example where System P is not satisfied is existential inference, where a query is valid if it follows from one repair.

On the non-satisfaction of the rational monotony (RM) property, the result is expected if we draw a parallel with standard possibilistic propositional logic (LP) and with the properties of non-monotonic relations. Indeed, there is a representation theorem (KLM [24]) which shows that any non-monotonic relation which satisfies System P and RM is necessarily representable by a total order on the set of interpretations of propositional logic.

On the Extension to Richer Languages: From a semantic point of view, the definitions of Cπ-repair given in Sect. 3, that have been established within

the framework of DL-Lite$_\mathcal{R}$, remain valid for richer languages (provided that the notion of deductive closure of an ABox with respect to a TBox can be defined). Indeed, the general process given in Fig. 2 is not proper to DL-Lite$_\mathcal{R}$ and easily applies to richer languages (e.g. Existential Rules). The challenge here is at the computational level, since we need first to find an equivalent characterization (like we did in this paper using support/dominance) and then show that it is tractable. For instance, in many description logics where conflicts may be composed of any number of assertions (unlike DL-Lite where conflicts consist of at most two assertions [16]), the extension of the support/dominance characterization is possible. However, even if the conflict set is computed in polynomial time, the size of this set itself can be exponential w.r.t. the size of the ABox. In this case, tractability cannot be preserved.

Furthermore, the main idea behind query-answering from inconsistent partially ordered knowledge bases is to extend the partial order into the set of its compatible total orders, then to apply a repair semantics to each one of them. The strategy we used in our approach (based on the possibilistic version of DL-Lite$_\mathcal{R}$) yields a single repair for each total order. However, in the general case, using a different strategy, each total order may return multiple repairs. Hence, a query needs to follow from all the repairs for all the compatible total orders.

On the Extension to Non-repair Based Semantics: We end this paper with a brief discussion on the applicability of inconsistency-tolerant semantics that are not directly based on repairs, such as paraconsistent multi-valued description logics, on partially ordered ABoxes. Let us first specify that an advantage of our approach is that once the possibilistic repair Cπ-repair is calculated, query-answering is done in a standard way. Within multi-valued semantics, the ABox remains unchanged, but the query-answering mechanisms need to be adapted and this can potentially generate an additional computational cost. Besides, from a semantic point of view, it is possible to redefine this work with multi-valued semantics. This can be done by first selecting a multi-valued semantics of DL-Lite$_\mathcal{R}$ (for example the 4-valued semantics given in [30]). The next step consists in extending it to the possibilistic framework with a totally ordered ABox. This requires an adaptation of the existing work (for flat ABox) to define preferred 4-valued canonical models. The last step consists in taking all the extensions of the total orders and defining the 4-valued canonical models of the partial ABox as the union of the preferred 4-valued canonical model of each total ABox extension. However, having an equivalent characterization (without generating all the extensions of the partial order) to the one given in this paper (Propositions 1 and 2), is not obvious to achieve.

References

1. Artale, A., Calvanese, D., Kontchakov, R., Zakharyaschev, M.: The DL-Lite family and relations. J. Artif. Intell. Res. (JAIR) **36**, 1–69 (2009). https://doi.org/10.1613/jair.2820
2. Baget, J., et al.: A general modifier-based framework for inconsistency-tolerant query answering. In: Principles of Knowledge Representation and Reasoning (KR), Cape Town, South Africa, pp. 513–516 (2016)
3. Baget, J.F., et al.: Inconsistency-tolerant query answering: rationality properties and computational complexity analysis. In: Michael, L., Kakas, A. (eds.) JELIA 2016. LNCS (LNAI), vol. 10021, pp. 64–80. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-48758-8_5
4. Banerjee, M., Dubois, D., Godo, L., Prade, H.: On the relation between possibilistic logic and modal logics of belief and knowledge. J. Appl. Non-Classical Logics **27**(3–4), 206–224 (2017)
5. Belabbes, S., Benferhat, S.: Computing a possibility theory repair for partially preordered inconsistent ontologies. IEEE Trans. Fuzzy Syst. **30**, 3237–3246 (2021)
6. Belabbes, S., Benferhat, S., Chomicki, J.: Handling inconsistency in partially pre-ordered ontologies: the Elect method. J. Log. Comput. **31**(5), 1356–1388 (2021)
7. Benferhat, S., Bouraoui, Z., Tabia, K.: 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, pp. 1450–1456 (2015)
8. Benferhat, S., Garcia, L.: Handling locally stratified inconsistent knowledge bases. Stud. Logica. **70**(1), 77–104 (2002)
9. Bienvenu, M.: On the complexity of consistent query answering in the presence of simple ontologies. In: Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence, Toronto, Ontario, Canada (2012)
10. Bienvenu, M., Bourgaux, C.: Inconsistency-tolerant querying of description logic knowledge bases. In: Pan, J.Z., et al. (eds.) Reasoning Web 2016. LNCS, vol. 9885, pp. 156–202. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-49493-7_5
11. Bienvenu, M., Bourgaux, C.: Querying and repairing inconsistent prioritized knowledge bases: complexity analysis and links with abstract argumentation. In: Principles of Knowledge Representation and Reasoning (KR), Virtual Event, pp. 141–151 (2020)
12. Bienvenu, M., Bourgaux, C., Goasdoué, F.: Querying inconsistent description logic knowledge bases under preferred repair semantics. In: AAAI, pp. 996–1002 (2014)
13. Brewka, G.: Preferred subtheories: an extended logical framework for default reasoning. In: IJCAI, Detroit, USA, pp. 1043–1048 (1989)
14. Britz, K., Casini, G., Meyer, T., Moodley, K., Sattler, U., Varzinczak, I.: Principles of KLM-style defeasible description logics. ACM Trans. Comput. Log. **22**(1), 1:1–1:46 (2021)
15. Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M., Rosati, R.: Tractable reasoning and efficient query answering in description logics: the DL-Lite family. J. Autom. Reason. **39**(3), 385–429 (2007)

16. Calvanese, D., Kharlamov, E., Nutt, W., Zheleznyakov, D.: Evolution of *DL-Lite* knowledge bases. In: Patel-Schneider, P.F., et al. (eds.) ISWC 2010. LNCS, vol. 6496, pp. 112–128. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-17746-0_8

17. Dempster, A.P.: Upper and lower probabilities induced by a multivalued mapping. Ann. Math. Stat. **38**, 325–339 (1967)

18. Dubois, D., Prade, H.: Possibilistic logic - an overview. Comput. Logic **9**, 197–255 (2014)

19. Dubois, D., Prade, H.: A crash course on generalized possibilistic logic. In: Ciucci, D., Pasi, G., Vantaggi, B. (eds.) SUM 2018. LNCS (LNAI), vol. 11142, pp. 3–17. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00461-3_1

20. Dubois, D., Prade, H., Schockaert, S.: Generalized possibilistic logic: foundations and applications to qualitative reasoning about uncertainty. Artif. Intell. J. **252**, 139–174 (2017)

21. Everett, L., Morris, E., Meyer, T.: Explanation for KLM-style defeasible reasoning. In: Jembere, E., Gerber, A.J., Viriri, S., Pillay, A. (eds.) SACAIR 2021. CCIS, vol. 1551, pp. 192–207. Springer, Cham (2022). https://doi.org/10.1007/978-3-030-95070-5_13

22. Fagin, R., Halpern, J.Y., Moses, Y., Vardi, M.Y.: Reasoning About Knowledge. MIT Press, Cambridge (2003)

23. Finger, M., Godo, L., Prade, H., Qi, G.: Advances in weighted logics for artificial intelligence. Int. J. Approximate Reasoning **88**, 385–386 (2017)

24. Kraus, S., Lehmann, D., Magidor, M.: Nonmonotonic reasoning, preferential models and cumulative logics. Artif. Intell. J. **44**(1–2), 167–207 (1990)

25. Lembo, D., Lenzerini, M., Rosati, R., Ruzzi, M., Savo, D.F.: Inconsistency-tolerant semantics for description logics. In: Hitzler, P., Lukasiewicz, T. (eds.) RR 2010. LNCS, vol. 6333, pp. 103–117. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-15918-3_9

26. Nebel, B.: Base revision operations and schemes: semantics, representation and complexity. In: European Conference on Artificial Intelligence, pp. 341–345 (1994)

27. Poggi, A., Lembo, D., Calvanese, D., Giacomo, G.D., Lenzerini, M., Rosati, R.: Linking data to ontologies. J. Data Semant. **10**, 133–173 (2008)

28. Shafer, G.: A Mathematical Theory of Evidence. Princeton University Press, Princeton (1976)

29. Spohn, W.: The Laws of Belief - Ranking Theory and Its Philosophical Applications. Oxford University Press, Oxford (2014)

30. Zhou, L., Huang, H., Qi, G., Ma, Y., Huang, Z., Qu, Y.: Paraconsistent query answering over DL-Lite ontologies. Web Intell. Agent Syst. Int. J. **10**(1), 19–31 (2012)