

Using ADFs for Inconsistency-Tolerant Query Answering with Existential Rules

Atefeh Keshavarzi Zafarghandi*, Patrick Koopmann

Vrije Universiteit Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands

Abstract

We present a new reduction of inconsistency-tolerant query answering to acceptance in ADFs. In particular, we consider knowledge bases (KBs) that use existential rules, and consider common inconsistency-tolerant semantics based on maximal consistent subsets. While reductions of inconsistency-tolerant reasoning to argumentation frameworks have been considered before, we aim to obtain a reduction that reflects the inference structure of the KB on a fine-grained level, so that they can be used to explain query answers on the level of individual inference steps. In particular, in our ADFs, every node corresponds to a fact derived in the *chase*, and acceptance conditions are used to relate facts using inference rules and integrity constraints. We show that our reduction satisfies rationality postulates, and observe that common semantics of ADFs fail to fully reproduce inconsistency-tolerant query answering with our reduction. We introduce a new semantics as refinement of the preferred semantics, which solves this problem, and analyze the computational complexity of this new semantics in the general and in our case.

Keywords

Abstract Dialectical Frameworks, Inconsistency-Tolerant Reasoning, Existential Rules

1. Introduction

One of the main advantages of symbolic AI is explainability. However, while most symbolic AI systems are explainable in theory, in practice, providing human understandable explanations remains a challenge that recently received a lot of attention [1, 2, 3]. In this paper, we focus on symbolic AI systems that are based on rule-languages, more particular extensions of datalog known as existential rules or tuple-generating dependencies (TGDs) [4]. Such rules allow to formulate domain knowledge, which can then be used in combination with a database to derive new facts to be queried by the user. The extension of a database with such a set of rules is then called a *knowledge base* (KB). This framework is also relevant in the ontology-based data access (OBDA) paradigm, since many ontologies can be translated into rules in such languages [5]. The problem we are concerned with is how to explain why a certain query answer can be derived from the data.

While earlier works on explainability focus on providing subsets of the data and/or rule sets as explanations [6, 7], more recent work also incorporates individual inference steps in the explanations [2, 8], to provide explanations in the form of proofs, which allows users to understand not only which facts are relevant for a query answer, but also why those facts are relevant. However, these work rely on classical semantics, and thus under the assumption that the data is consistent with the background knowledge. In realistic scenarios, this assumption cannot always be made, which is why different *inconsistency-tolerant semantics* have been introduced to allow to provide meaningful answers also for inconsistent datasets. There are different approaches to explain query answers under inconsistency-tolerant semantics. However, they all rely on directly linking sets of facts from the database directly to the conclusion, and thus are not able to provide more detailed explanations comparable to those of proofs. For instance, [9] proposes to use *argumentation*

frameworks (AFs) to represent inconsistency-tolerant reasoning, where each *argument* has itself a complex structure consisting of a *support* (facts from the database) and *conclusions* (facts derivable from the support), but the relation between those two sets is not explained.

Motivated by this short-coming, we introduce a new translation of inconsistent KBs into Abstract Dialectical Frameworks (ADFs) [10]. Similar to AFs, (ADFs) can be visualized using directed graphs. However, while in AFs, nodes represent abstract arguments, and edges depict conflicts (attack relations), in ADFs nodes can represent arguments, propositions, or even statements. The edges between them can signify attack, support, or even a combination of both, depending on the specific context, which is formalised using *acceptance conditions* in the form of propositional formulas. Our translation is based on the *chase*, which is a common formalism to capture inferences of rule-based KBs. We represent every fact that is derived in the chase as node in the ADF, and use acceptance conditions to encode the role of this fact in the computation of the chase (which facts does it depend on, how do other facts depend on it), as well as possible conflicts with other derived facts. Since those acceptance conditions are transparent and directly related to the rules in the program, we believe that those ADFs can be used to provide detailed explanations of query answers under different inconsistency-tolerant semantics in a similar way to proof-based explanations.

We show that the translated ADFs satisfy desirable rationality postulates for such a translation, when using the well-established *admissible* and *preferred* semantics. However, it turns out that the preferred semantics does not capture the usual inconsistency-tolerant semantics of KBs [9], which is why we introduce a new semantics, the *min- S' -max preferred interpretation*, for which we show that it does replicate standard inconsistency-tolerant semantics. We analyze the computational complexity of deciding acceptance in ADFs under *min- S' -max preferred interpretation*, and then look at the special case where the ADF is the result of our translation from an inconsistent KB. Thanks to the syntactical shape of our acceptance conditions, we are able to show that the complexity of acceptance under *min- S' -max preferred interpretation* drops by one polynomial, which allows us to (re-)prove complexity results for inconsistency-tolerant query answering.

22nd International Workshop on Nonmonotonic Reasoning, November 2-4, 2024, Hanoi, Vietnam

✉ a.keshavarzi.zafarghandi@vu.nl (A. Keshavarzi Zafarghandi);

p.k.koopmann@vu.nl (P. Koopmann)

🆔 0000-0002-5806-1012 (A. Keshavarzi Zafarghandi);

0000-0001-5999-2583 (P. Koopmann)

© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



2. Related Work

Since meaningful query answers cannot be obtained from inconsistent KBs under classical semantics based on first order logics, recent research has explored various approaches to address this challenge. 1. *Inconsistency-Tolerant Semantics*: This approach focuses on developing semantics for query answering that can handle inconsistencies in the KB. Some works define repairs, which are subsets of the KB that are consistent, and then use these repairs to determine the answer to a query [11, 7, 12] 2. *Argumentation Formalisms*: Here, the KB is transformed into an argumentation framework. The framework then identifies justification of arguments [13, 14, 9].

While query answering has long been a cornerstone of database research (e.g., [15, 16, 17]), recent research started also looking closer on possible ways of explaining query answers: [6] introduces a proof-theoretic approach for this in the case the KB uses an ontology in DL-Lite. [2] considers programs with existential rules, and provides explanations that combine proofs and universal models based on the chase. [7] tackles explaining query answers in the context of inconsistency-tolerant query answering for existential rules, focusing on three popular inconsistency-tolerant semantics. While the focus of [11] is on querying inconsistent description logic KBs, the paper also explores explaining why a tuple is a (non-)answer under different semantics.

A number of works link argumentation and non-monotonic reasoning. The process of converting KBs into AFs has been a well-studied area (e.g., [14, 18, 19, 9]). However, with the exception of [9], those works consider propositional defeasible and default logics, while we are looking at first-order theories (in particular: existential rules), and inconsistency-tolerant semantics. Converting KBs into AFs involves extracting arguments and their relationships (support, attack) from the KB, creating an abstract representation of arguments and their interconnections. This abstraction offers several advantages: 1. *Generality*: The framework can be applied across diverse domains like legal reasoning, decision making, and diagnosis. 2. *Content-Independent Evaluation*: Argument evaluation focuses solely on structural relationships, independent of specific content details. This evaluation is guided by semantics, as defined in [20, 21]. 3. *Explainability*: Semantics can be interpreted dialectically, allowing explanations for inferred conclusions (e.g., [22, 23]).

However, abstraction also has limitations. While abstract formalisms and their semantics aim to identify a conflict-free set of arguments, these arguments might still not draw rational conclusions when mapped back to the original KB [24, 25]. Thus, a crucial question in literature is: how do we construct an abstract formalism from a KB to draw rational conclusions? [14, 24, 19, 25, 13] aims to retain the theoretical and computational benefits of abstract argumentation formalisms while addressing the instantiation challenges. [18] proposes rationality postulates as a set of criteria for evaluating argumentation formalisms. Their work emphasizes the importance of properties like closure and direct/indirect consistency for ensuring well-founded reasoning processes.

Building on their previous work ([26, 27]), [19] explores the characteristics an argumentation framework should possess to fulfill rationality postulates. These postulates act as a set of constraints that guide the selection of arguments within abstract formalisms, ensuring that the reasoning pro-

cess remains consistent and avoids conflicts present in the underlying KB.

To leverage the advantages of abstract argumentation formalisms—namely, systematically dealing with inconsistencies and using the structure of the argumentation graph to enhance transparency—the first task of this work is to transform a KB to an abstract argumentation formalism. The second task is to investigate whether the induced argumentation formalism from the given KB satisfies rationality postulates.

While several works have instantiated KBs into abstract argumentation frameworks (AFs) and show that the resulting framework satisfies rationality postulates (e.g., [24, 18, 14]), [13] presented an alternative approach by instantiating KBs in expressive generalization of AFs, i.e., abstract dialectical frameworks (ADFs) [28, 29].

While [13] proposes a method for translating defeasible theory bases (containing both defeasible and strict rules) into Abstract Dialectical Frameworks (ADFs), in our work we focus on KBs containing only strict rules.

Since ADFs allow nodes to represent statements, each fact from a strict KB can be directly translated into an abstract node within the ADF. This simplified approach enables the construction of the associated ADF from a KB in polynomial time with respect to the size of a chase of a KB.

Furthermore, in our paper we show that the resulting ADF satisfies the rationality postulates under preferred semantics. Additionally, a new type of semantics for ADFs is introduced. This new semantics is shown to be equivalent to the chase procedure applied to a repair of the original KB.

3. Preliminaries

We assume the reader is familiar with basics of first-order logics. Given sets N and M , we use $N \setminus M$ as alternative notation for $N \cap M$.

3.1. Abstract Dialectical Frameworks

We summarize key concepts of ADFs [29, 28].

Definition 1. An abstract dialectical framework (ADF) is a tuple $D = (S, L, C)$ where:

1. S is a set of statements (arguments, positions);
2. $L \subseteq S \times S$ is a set of links among statements;
3. $C = \{\varphi_s\}_{s \in S}$ contains, for each $s \in S$, a propositional formula φ_s over

$$\text{par}(s) = \{b \in S \mid (b, s) \in L\}.$$

An ADF can be represented by a graph in which nodes indicate arguments/statements and links show the relations between them. Intuitively, the acceptance condition of each statement clarifies under which conditions the statement can be accepted.

An *interpretation* v (for D) is a function $v : S \mapsto \{\mathbf{t}, \mathbf{f}, \mathbf{u}\}$ that maps statements to one of the three truth values *true* (\mathbf{t}), *false* (\mathbf{f}), and *undecided* (\mathbf{u}). v is called *two-valued* if for every $s \in S$, $v(s) \in \{\mathbf{t}, \mathbf{f}\}$. The *trivial* interpretation $v_{\mathbf{u}}$ satisfies $v_{\mathbf{u}}(s) = \mathbf{u}$ for all $s \in S$. For brevity, we sometimes identify v with the corresponding set of literals

$$\{s_i \mid v(s_i) = \mathbf{t}\} \cup \{\neg s_i \mid v(s_i) = \mathbf{f}\}.$$

For instance, $v = \{a \rightarrow \mathbf{f}, b \rightarrow \mathbf{t}, c \rightarrow \mathbf{u}\}$ corresponds to $\{-a, b\}$. Furthermore, we set $v^{\mathbf{t}} = \{s \mid v(s) = \mathbf{t}\}$. We lift interpretations v to propositional formulas $v(\varphi) = \varphi^v$ over S , where φ^v is obtained from φ by replacing every $s \in S$ by \top if $v(s) = \mathbf{t}$, by \perp if $v(s) = \perp$, and keeping it if $v(s) = \mathbf{u}$. Given a statement $s \in S$, s is called *acceptable* w.r.t. v if φ_s^v is irrefutable (a tautology). The semantics are defined via the *characteristic operator* Γ_D .

Definition 2. Let D be an ADF and let v be an interpretation of D . The characteristic operator Γ_D is defined by $\Gamma_D(v) = v'$, where for each $s \in S$, $v'(s) = \mathbf{t}$ if φ_s^v is irrefutable, $v'(s) = \mathbf{f}$ if φ_s^v is unsatisfiable, and $v'(s) = \mathbf{u}$, otherwise.

The *information ordering* $<_i$ on truth values is the smallest ordering satisfying $\mathbf{u} <_i \mathbf{t}$ and $\mathbf{u} <_i \mathbf{f}$. \leq_i is the reflexive closure of $<_i$. We extend \leq_i to interpretations by setting $w \leq_i v$ if $w(s) \leq_i v(s)$ for each $s \in S$.

A *semantics* σ for ADFs now assigns to each ADF a set of interpretations. Most types of semantics for ADFs are based on the concept of admissibility. An interpretation v for a given ADF D is called *admissible* iff $v \leq_i \Gamma_D(v)$; it is *preferred* iff v is \leq_i -maximal admissible; it is *complete* iff $v = \Gamma_D(v)$; it is the *grounded* interpretation of D iff v is the least fixed point of Γ_D ; it is a (*two-valued*) *model* iff v is two-valued and $\Gamma_D(v) = v$; it is *stable* iff v is a two-valued model of D and $v^{\mathbf{t}} = w^{\mathbf{t}}$, where w is the grounded interpretation of the *stb*-reduct $D^v = (S^v, L^v, C^v)$, where $S^v = v^{\mathbf{t}}$, $L^v = L \cap (S^v \times S^v)$, and $\varphi_s[p/\perp : v(p) = \mathbf{f}]$ for each $s \in S^v$.

The set of all σ interpretations for an ADF D is denoted by $\sigma(D)$, where $\sigma \in \{\text{adm}, \text{prf}, \text{com}, \text{grd}, \text{mod}, \text{stb}\}$ abbreviates the different semantics in the obvious manner.

Definition 3. Let σ a semantics for ADFs, $D = (S, L, C)$ be an ADF, $s \in S$ and v and interpretation.

1. The verification problem, denoted by $\text{Ver}_\sigma(v, D)$, asks whether $v \in \sigma(D)$.
2. The credulous acceptance problem, denoted by $\text{Cred}_\sigma(s, D)$, asks if for some $v \in \sigma(D)$, $v(s) = \mathbf{t}$.
3. The skeptical acceptance problem, denoted by $\text{Skept}_\sigma(s, D)$, asks if for each $v \in \sigma(D)$, $v(s) = \mathbf{t}$.

3.2. Existential Rules

An *atom* is an atomic formula $A(\vec{t})$ for some predicate A , which is *ground* if \vec{t} contains only constants. An *instance* I is a set of first-order *ground atoms*, which is called a *database* if it is finite. For convenience, we identify first-order interpretations with instances. An *existential rule* is a first-order formula of the form $\forall \vec{x}, \vec{y} \phi(\vec{x}, \vec{y}) \rightarrow \exists \vec{z} \psi(\vec{y}, \vec{z})$, and an *integrity constraint* is a formula of the form $\forall \vec{x}, \vec{y} \phi(\vec{x}, \vec{y}) \rightarrow \perp$, where $\psi(\vec{y}, \vec{z})$ are possibly empty conjunctions of atoms. We call $\phi(\vec{x}, \vec{y})$ the *body*, and $\exists \vec{z} \psi(\vec{y}, \vec{z})$ the *head* of the rule, and treat those conjunctions as *sets*, i.e. the order of conjuncts is not relevant. Existential rules and integrity constraints are collectively called *rules*, and for convenience, we leave out the universal quantification when writing them. A *program* is a set \mathcal{P} of rules, and a *knowledge base* (KB) is a tuple $\mathcal{K} = \langle \mathcal{D}, \mathcal{P} \rangle$ with \mathcal{D} a database and \mathcal{P} a program, which we identify with the first-order theory $\mathcal{D} \cup \mathcal{P}$.

A *conjunctive query* (CQ) is a first-order formula of the form $q(\vec{x}) = \exists \vec{y} \phi(\vec{x}, \vec{y})$, where $\phi(\vec{x}, \vec{y})$ is a conjunction of

	DL Syntax	Corresponding Rule
(A1)	$A \sqsubseteq B$	$A(x) \rightarrow B(x)$
(A2)	$A_1 \sqcap A_2 \sqsubseteq B$	$A_1(x) \wedge A_2(x) \rightarrow B(x)$
(A3)	$A \sqcap B \sqsubseteq \perp$	$A(x) \wedge B(x) \rightarrow \perp$
(A4)	$A \sqsubseteq \exists R.B$	$A(x) \rightarrow \exists x.R(x, y) \wedge B(y)$
(A5)	$A \sqsubseteq \forall R.B$	$A(x) \wedge R(x, y) \rightarrow B(y)$
(A6)	$A \sqsubseteq \exists R.\text{Self}$	$A(x) \rightarrow R(x, x)$
(A7)	$R \sqsubseteq S$	$R(x, y) \rightarrow S(x, y)$
(A8)	$R \circ S \sqsubseteq T$	$R(x, y) \wedge S(y, z) \rightarrow T(x, z)$

Table 1

Axioms of Horn- $\mathcal{SR}\mathcal{I}$, where A, A_1, A_2, B are unary predicates, R, S are binary predicates r or their *inverse* r^- , where we identify $r^-(x, y)$ with $r(y, x)$.

atoms over \vec{x}, \vec{y} and the set of constants. An *answer* to this query is a vector \vec{c} of constants s.t. $\mathcal{K} \models q(\vec{c})$.

Query entailment for KBs can be characterized using the *chase*, and in this paper, it is convenient to use the so-called Skolem chase for this. To define the Skolem chase, we first define the *Skolemization* \mathcal{P}_{sk} of \mathcal{P} , which is obtained from \mathcal{P} by replacing every existentially quantified variable z_i in an existential rule $\phi(\vec{x}, \vec{y}) \rightarrow \exists \vec{z} \psi(\vec{y}, \vec{z})$ by the term $f(\vec{x}, \vec{y})$, where f is a function symbol that is unique to that rule and the variable z_i . The *grounding* \mathcal{P}_D^{ch} of \mathcal{P} wrt \mathcal{D} contains all ground rules that can be obtained by replacing the variables in \mathcal{P}_{sk} with terms obtained using the constants and function symbols in $\mathcal{P}_{sk} \cup \mathcal{D}$. Note that $\mathcal{D} \cup \mathcal{P}_D^{ch}$ is essentially a possibly infinite set of propositional Horn formulas.

A ground existential rule can be *applied* on an instance I if its body occurs in I , and the result of this application is then obtained by adding its head to I . The *Skolem chase*, denoted by $\text{chase}(\mathcal{K})$, is now the fixpoint of a fair application of the existential rules in \mathcal{P}_D^{ch} on \mathcal{D} . Here, *fair* means that every applicable rule is eventually applied.

It is well known that \mathcal{K} is consistent iff $I = \text{chase}(\mathcal{K})$ does not invalidate any integrity constraint in \mathcal{P} . In this case, for every CQ, the answers over \mathcal{K} are exactly the answers over $\text{chase}(\mathcal{K})$.

In the context of this paper, we look at fragments of the full language of existential rules. Many description logics of the Horn-family can be translated into rules (see Table 1). We focus on programs for which $\text{chase}(\mathcal{K})$ is finite, which is for instance the case if

- \mathcal{P} is *datalog*, i.e. it contains neither existentially quantified variables nor function symbols, or
- \mathcal{P} is *cycle-restricted* (see [30] for a formal definition).

3.3. Inconsistency-Tolerant Reasoning

Since no meaningful query answers can be obtained under the usual first-order semantics if $\mathcal{K} \models \perp$, several inconsistency-tolerant semantics based on ABox repairs have been proposed in the literature [31].

Definition 4 (ABox repair). An ABox repair for a KB $\mathcal{K} = \langle \mathcal{D}, \mathcal{P} \rangle$ is a \subseteq -maximal $\mathcal{D}' \subseteq \mathcal{D}$ s.t. $\langle \mathcal{D}', \mathcal{P} \rangle \not\models \perp$. We use $\text{Repairs}(\mathcal{D}, \mathcal{P}) = \text{Repairs}(\mathcal{K})$ to denote the set of all such repairs.

In this paper, we focus on the following semantics for inconsistency-tolerant query answering.

Definition 5. Given a KB $\mathcal{K} = \langle \mathcal{D}, \mathcal{P} \rangle$ a query $q(\vec{x})$ and a vector \vec{c} of constants the length of \vec{x} ,




Figure 1: The acceptance condition of B in the induced ADF $D(\mathcal{K})$ is $\varphi_B = \varphi_B^1 \wedge \varphi_B^2 \wedge \varphi_B^3$, such that $\varphi_B^1 = \bigwedge_{i=1}^n C_i \rightarrow D$, $\varphi_B^2 = \bigwedge_{i=1}^n A_i$, and $\varphi_B^3 = \neg E$, where $\bigwedge_{i=1}^n A_i \rightarrow B \in \mathcal{P}_D^{ch}$, $B \wedge \bigwedge_{i=1}^n C_i \rightarrow D \in \mathcal{P}_D^{ch}$, and $B \wedge E \rightarrow \perp \in \mathcal{P}_D^{ch}$.

1. \vec{c} is an answer under AR semantics, in symbols $\mathcal{K} \models_{AR} q(\vec{c})$, if for every repair $\mathcal{A}' \in \text{Repairs}(\mathcal{K})$, $\langle \mathcal{P}, \mathcal{A}' \rangle \models q(\vec{c})$.
2. \vec{c} is an answer under brave semantics, denoted $\mathcal{K} \models_{brave} q(\vec{c})$, if for some repair $\mathcal{A}' \in \text{Repairs}(\mathcal{K})$, $\langle \mathcal{P}, \mathcal{A}' \rangle \models q(\vec{c})$.

4. Induced ADF from a given KB

Our idea is to represent atoms in the chase directly as statements in the ADF, and use links to capture rule inferences as well as conflicts detected through integrity constraints. This way, understanding why a query is entailed under a certain semantics or not can be exhibited by navigating through the graph and analyzing the propagation of facts and conflicts.

We show that our translation satisfies common rationality postulates, but also that standard semantics for ADFs are not able to reproduce usual semantics for inconsistency-tolerant reasoning with our formalization. We thus introduce a new semantics for ADFs called min- S' -max preferred semantics (Definition 8). In particular, we will be able to show that the extensions wrt that semantics correspond to chases of repairs.

We focus on queries of a single atom, as more complex queries can be reduced to them: for a given query $q(\vec{x}) \leftarrow \exists \vec{y}. \phi(\vec{x}, \vec{y})$, we would simply add the rule $\phi(\vec{x}, \vec{y}) \rightarrow Q(\vec{x})$, where Q is a fresh predicate. Another assumption we make w.l.o.g. is that integrity constraints always contain exactly two atoms. Recall that \mathcal{P}_D^{ch} refers to the grounding of \mathcal{P} .

Definition 6 (Induced ADF). Let $\mathcal{K} = \langle \mathcal{D}, \mathcal{P} \rangle$ be KB. Its induced ADF $D(\mathcal{K}) = (S_{\mathcal{K}}, L_{\mathcal{K}}, C_{\mathcal{K}})$ is defined as follows:

1. $S_{\mathcal{K}} = \text{chase}(\mathcal{K})$
2. for each $B \in S_{\mathcal{K}}$, $C_{\mathcal{K}}$ contains the acceptance condition φ_B of B defined as $\varphi_B = \varphi_B^1 \wedge \varphi_B^2$ if $B \in \mathcal{D}$, and otherwise as $\varphi_B = \varphi_B^1 \wedge \varphi_B^2 \wedge \varphi_B^3$, where,

$$\varphi_B^1 = \bigwedge \left\{ \bigwedge_{i=1}^n C_i \rightarrow D \mid B \wedge \bigwedge_{i=1}^n C_i \rightarrow D \in \mathcal{P}_D^{ch}, C_i \in S_{\mathcal{K}}, D \in S_{\mathcal{K}} \right\}$$

$$\varphi_B^2 = \bigvee \left\{ \bigwedge_{i=1}^n A_i \mid \bigwedge_{i=1}^n A_i \rightarrow B \in \mathcal{P}_D^{ch}, A_i \in S_{\mathcal{K}} \right\}$$

$$\varphi_B^3 = \bigwedge \left\{ \neg E \mid B \wedge E \rightarrow \perp \in \mathcal{P}_D^{ch}, E \in S_{\mathcal{K}} \right\}$$

3. $L_{\mathcal{K}}$ contains $\langle A, B \rangle$ for every A and B s.t. A occurs in φ_B

Figure 1 shows a partial illustration of the induced ADF.

The intuition behind φ_B^3 of B is as follows: if there is an integrity constraint $B \wedge E \rightarrow \perp$, and $B, E \in \text{chase}(\mathcal{K})$, then B and E cannot be both in the chase of a repair.

φ_B^2 intuitively ensures that every accepted statement is justified through inferences in the chase, starting from the facts in a repair. If B is a head of some rule, B is in the chase of a repair \mathcal{A}' iff all the premises (body) of at least one of the rules are in the chase of \mathcal{A}' .

Finally, the intuition behind φ_B^1 is as follows: if B occurs in the body of a rule and all atoms in the body are satisfied, then the head D has to be satisfied as well. If the acceptance condition of D makes it impossible for D to be satisfied, e.g. because D conflicts with another atom that we assigned to true, then B cannot be accepted either, since otherwise a rule in \mathcal{P} would not be satisfied in the interpretation.

We demonstrate the soundness of our transformation. To achieve this, we adopt a two-step approach. First, we establish the rationality of the constructed ADF. We do this by verifying that the induced ADF satisfies the well-known rational postulates for argumentation formalisms [24, 32], under preferred and min- S' -max preferred semantics, as presented in Section 4.1.

Second, we investigate the correspondence between interpretations of our ADFs and KB repairs. We illustrate with examples why the current semantics of ADFs is insufficient for capturing exactly the set of repairs and their chases. We then define a new semantics based on a subclass of preferred interpretations, and show that this semantics does allow us to reproduce query entailment under standard inconsistency-tolerant semantics using credulous and skeptical acceptance.

Before going to the formal proofs, we illustrate the transformation process with some examples. We first look at the consistent case.

Example 1. Let $\mathcal{K} = \langle \mathcal{D}, \mathcal{P} \rangle$ be KB such that $\mathcal{D} = \{A(a)\}$ and $\mathcal{P} = \{A(x) \rightarrow B(x)\}$. Then the chase of \mathcal{K} is $\{A(a), B(a)\}$. The set of statements in the induced ADF $D(\mathcal{K})$ is $S_{\mathcal{K}} = \{A(a), B(a)\}$. The acceptance conditions are

$$\varphi_{A(a)} : \top \rightarrow B(a) \quad \varphi_{B(a)} : A(a)$$

$D(\mathcal{K})$ has two preferred interpretations: $v_1 = \{A(a), B(a)\}$ and $v_2 = \{\neg A(a), \neg B(a)\}$, and only v_2 is a stable model of $D(\mathcal{K})$.

Intuitively, a repair of a KB is a maximal consistent set of facts from the database. Similarly, a preferred interpretation of an ADF contains the maximal information (w.r.t. \leq_i -ordering) about the acceptance of statements. In Example 1, while interpretation v_1 accepts both A and B , v_2 rejects both. The repair of \mathcal{K} is $v_1 \upharpoonright_{\mathcal{D}}$, and $v_1 \upharpoonright_{\mathcal{D}}$ is the chase of this repair. Thus, although preferred interpretations contain the maximum information about the acceptance of the statements, since all statements that support each other in a loop can be assigned to **f** together in a preferred interpretation, the set of preferred interpretations does not coincide with the set of repairs of the associated KB. Moreover, this example demonstrates a distinction between the set of stable interpretations of the induced ADF and the repairs of the given KB. Example 2 contains an integrity constraint.

Example 2. Let $\mathcal{K} = \langle \mathcal{D}, \mathcal{P} \rangle$ be KB such that $\mathcal{D} = \{A(a), C(a)\}$ and $\mathcal{P} = \{A(x) \rightarrow B(x), C(x) \rightarrow D(x), B(x) \wedge D(x) \rightarrow \perp\}$. Then the set of statements

in $D(\mathcal{K})$ is $S_{\mathcal{K}} = \{A(a), B(a), C(a), D(a)\}$ and the acceptance conditions are

$$\begin{aligned}\varphi_{A(a)} &: \top \rightarrow B(a) & \varphi_{C(a)} &: \top \rightarrow D(a) \\ \varphi_{B(a)} &: A(a) \wedge \neg D(a) & \varphi_{D(a)} &: C(a) \wedge \neg B(a)\end{aligned}$$

ADF $D(\mathcal{K})$ has two preferred interpretations:

$$\begin{aligned}v_1 &= \{A(a), B(a), \neg C(a), \neg D(a)\} \\ v_2 &= \{\neg A(a), \neg B(a), C(a), D(a)\}.\end{aligned}$$

Observe that $v_1^{\mathbf{t}} \not\subseteq v_2^{\mathbf{t}}$ and $v_2^{\mathbf{t}} \not\subseteq v_1^{\mathbf{t}}$. Both $v_1^{\mathbf{t}}|_{\mathcal{D}}$ and $v_2^{\mathbf{t}}|_{\mathcal{D}}$ are repairs for \mathcal{K} , furthermore $v_1^{\mathbf{t}}$ and $v_2^{\mathbf{t}}$ coincide with the chases of the repairs of \mathcal{K} .

While in Example 2, the set of preferred interpretations and the set of repairs coincide, we will later see that this is not always the case. Before we further explore the relationship between ADF semantics and the set of repairs for the KB in Section 4.2, we show that the induced ADF satisfies the rationality postulates defined in [24, 32].

4.1. Rationality Postulates

The first postulate requires that σ -interpretations should be closed under rule-applications, which in our case are captured by the chase.

Postulate 1 (Closure). *Given KB $\mathcal{K} = \langle \mathcal{D}, \mathcal{P} \rangle$, $D(\mathcal{K})$ satisfies closure for semantics σ iff for any σ -interpretation v , it holds that $v^{\mathbf{t}} = \text{chase}(v^{\mathbf{t}})$.*

Another crucial property of induced ADFs is consistency. The literature [18] defines two levels of consistency: *direct consistency* and *indirect consistency*.

Definition 7. *Given KB $\mathcal{K} = \langle \mathcal{D}, \mathcal{P} \rangle$, its induced ADF $D_{\mathcal{K}}$, we call an interpretation v is consistent if for each $A \wedge B \rightarrow \perp \in \mathcal{P}_{\mathcal{D}}^{\text{ch}}$ s.t. $v(A) = \mathbf{t}$, we have $v(B) \neq \mathbf{t}$.*

Postulate 2 (Direct Consistency). *Given KB $\mathcal{K} = \langle \mathcal{D}, \mathcal{P} \rangle$ and its induced ADF $D(\mathcal{K})$, $D(\mathcal{K})$ satisfies direct consistency for semantics σ iff each σ -interpretation v is consistent.*

Postulate 3 (Indirect Consistency). *Given KB $\mathcal{K} = \langle \mathcal{D}, \mathcal{P} \rangle$ and its induced ADF $D(\mathcal{K})$. $D(\mathcal{K})$ satisfies indirect consistency for semantics σ iff for each σ -interpretation v , $Cl(v)$ is consistent.*

We show that the induced ADF defined in Definition 6 satisfy those rationality postulates under admissible semantics.

Theorem 1. *For every KB \mathcal{K} , $D(\mathcal{K})$ satisfies closure for admissible semantics.*

Proof. Let $\mathcal{K} = \langle \mathcal{D}, \mathcal{P} \rangle$. We show that for each admissible interpretation v of $D(\mathcal{K})$, $\text{chase}(v^{\mathbf{t}}) = v^{\mathbf{t}}$. $v^{\mathbf{t}} \subseteq \text{chase}(v^{\mathbf{t}})$ follows directly from the definition of the chase, so that we only need to show $\text{chase}(v^{\mathbf{t}}) \subseteq v^{\mathbf{t}}$. Let $\bigwedge_{i=1}^n A_i \rightarrow B \in \mathcal{P}_{\mathcal{D}}^{\text{ch}}$ be such that $A_i \in v^{\mathbf{t}}$ for each $1 \leq i \leq n$. We need to show that then, $B \in v^{\mathbf{t}}$. It follows from $\varphi_{A_i}^1$ in the acceptance condition of each A_i , and the fact that v is admissible, that is if $v(A_i) = \mathbf{t}$ for each A_i , then $(\bigwedge_{j \neq i} A_j \rightarrow B)^v \equiv \top$. Thus, since $v(A_i) = \mathbf{t}$ for each $1 \leq i \leq n$, we obtain $v(B) = \mathbf{t}$. Hence, $Cl(v) = v^{\mathbf{t}}$ for each admissible interpretation v . \square

Theorem 2. *For every KB \mathcal{K} , $D(\mathcal{K})$ satisfies direct consistency for admissible semantics.*

Proof. Toward a contradiction, assume $D(\mathcal{K})$ does not satisfy direct consistency for admissible semantics. Thus, there exists an admissible interpretation v of $D(\mathcal{K})$, and there exists $A \wedge B \rightarrow \perp \in \mathcal{P}_{\mathcal{D}}^{\text{ch}}$ s.t., $A, B \in v^{\mathbf{t}}$. Since, by Definition 6, $\neg B$ is in the acceptance condition of A and v is admissible, if $v(B) = \mathbf{t}$, we must have $v(A) = \mathbf{f}$. Thus, the assumption that $A, B \in v^{\mathbf{t}}$ cannot be true. \square

Corollary 3. *For every KB \mathcal{K} , $D(\mathcal{K})$ satisfies indirect consistency for admissible semantics.*

Proof. This follows from Theorem 1 and Theorem 2. \square

4.2. Minmax-preferred semantics for ADFs

Since a preferred interpretation of an ADF is a \leq_i -maximal admissible interpretation, the preferred semantics of ADFs can be a candidate for evaluating ABox repair semantics of the associated KB \mathcal{K} . In Example 2 (an inconsistent KB), each preferred interpretation of the induced ADF corresponds to a repair of the original KB, \mathcal{K} . However, in Example 1 (a consistent KB), only one of the preferred interpretations, which is not a stable model, is a repair. This demonstrates that stable models of the induced ADF do not have to correspond to repairs. Examples 1–2 indicate that, if $v^{\mathbf{t}}|_{\mathcal{D}}$ should coincide with a repair of \mathcal{K} , $v^{\mathbf{t}}|_{\mathcal{D}}$ needs to be subset-maximal among preferred interpretations. This is however not sufficient, as it is illustrated by Example 3.

Example 3. *Let $\mathcal{K} = \langle \mathcal{D}, \mathcal{P} \rangle$ be a KB such that $\mathcal{D} = \{A(a), B(a)\}$ and \mathcal{P} contains the following rules:*

$$\begin{aligned}A(x) \wedge B(x) &\rightarrow \perp & A(x) \wedge B(x) &\rightarrow C(x) \\ A(x) \wedge B(x) &\rightarrow D(x) & D(x) &\rightarrow C(x) & C(x) &\rightarrow D(x)\end{aligned}$$

We have as acceptance conditions:

$$\begin{aligned}\varphi_{A(a)} &: \neg B(a) \wedge (B(a) \rightarrow C(a)) \wedge (B(a) \rightarrow D(a)) \\ \varphi_{B(a)} &: \neg A(a) \wedge (A(a) \rightarrow C(a)) \wedge (A(a) \rightarrow D(a)) \\ \varphi_{C(a)} &: ((A(a) \wedge B(a)) \vee D(a)) \wedge D(a) \\ \varphi_{D(a)} &: ((A(a) \wedge B(a)) \vee C(a)) \wedge C(a)\end{aligned}$$

The following is a preferred interpretation that maximises $v^{\mathbf{t}}|_{\mathcal{D}}$, and indeed contains a repair:

$$v = \{A(a), \neg B(a), C(a), D(a)\}$$

However, the facts $C(a)$ and $D(a)$ are not entailed by that repair, and should not be considered for inconsistency tolerant query answers.

To reproduce inconsistency-tolerant semantics, our interpretations do not only need to capture the repairs, but also their entailments, as produced by the chase. Together, our examples illustrate what needs to be done to obtain a semantics that fully characterizes entailments of repairs. First, we need to *maximize* the set of statements that correspond to facts from the database. This aligns with the definition of repairs to be subset-maximal consistent subsets of the database. In addition, we have to *minimize* the set of derived statements, to reflect the minimality of the chase that contains only facts that can be deduced from the database.

In Definition 8, we define a new type of semantics for ADFs called min- S' -max preferred semantics, which uses

as parameter a set $S' \subseteq S$ of statements to be maximized. While our definition is generic to any set S' , we will later instantiate S' with \mathcal{D} , to address the first point made above. An interpretation v is considered $\text{min-}S'$ -max preferred if it is maximal not only with respect to the \leq_i ordering but also with respect to \subseteq concerning $S' \subseteq S$. That is, $v^{\mathbf{t}}|_{S'} = v^{\mathbf{t}} \cup S'$ has to be maximal among the preferred interpretations. Furthermore, $v^{\mathbf{t}}$ has to be *minimal* w.r.t. \subseteq relation among all chosen interpretation in the previous step.

Definition 8. Let $D = (S, L, C)$ be an ADF, and let $S' \subseteq S$. An interpretation v is $\text{min-}S'$ -max preferred interpretation (for D), in symbols $v \in \text{min-}S'$ -max-prf(D), iff $v \in \text{prf}(D)$ and

1. $v^{\mathbf{t}}|_{S'}$ is maximal among all preferred interpretations of D , i.e., if $w \in \text{prf}(D)$, then $v^{\mathbf{t}}|_{S'} \not\subseteq w^{\mathbf{t}}|_{S'}$,
2. $v^{\mathbf{t}}$ is minimal among all $w \in \text{prf}(D)$ s.t. $v^{\mathbf{t}}|_{S'} = w^{\mathbf{t}}|_{S'}$, i.e., $w^{\mathbf{t}} \not\subseteq v^{\mathbf{t}}$ for any such w .

Note that the order of the conditions in Definition 6 is crucial. First, we select a preferred interpretation v as a candidate. The first condition checks if $v^{\mathbf{t}}|_{S'}$ is maximal among all $w^{\mathbf{t}}|_{S'}$, where $w \in \text{prf}(D)$. Then, the second condition verifies if $v^{\mathbf{t}}$ is minimal among the interpretations chosen in the first step.

Returning to Example 3, we observe that for $S' = \mathcal{D}$, v is not $\text{min-}\mathcal{D}$ -max preferred, since $C(a)$ and $D(a)$ contradict the minimality criterion, but assigning those statements to \mathbf{f} , we obtain a $\text{min-}\mathcal{D}$ -max preferred interpretation, which corresponds to the chase of a repair as required.

Lemma 4. Let $D = (S, L, C)$ be an ADF, and let $S' \subseteq S$.

1. Every ADF has at least one $\text{min-}S'$ -max preferred interpretation w.r.t. S' .
2. Every $\text{min-}S'$ -max preferred interpretation is a preferred interpretation.
3. $\text{min-}S'$ -max preferred semantics differs from preferred semantics.
4. $\text{min-}S'$ -max preferred semantics differs from stable semantics.

Proof. Every ADF has at least one preferred interpretation. Among these preferred interpretations, there exists at least one interpretation v such that $v^{\mathbf{t}}|_{S'}$ is subset-maximal among all $w^{\mathbf{t}}|_{S'}$, where w is a preferred interpretation. Furthermore, within this set of interpretations, there exists at least an interpretation v such that $v^{\mathbf{t}}$ is minimal among all w , where $w^{\mathbf{t}} = v^{\mathbf{t}}$ and $w \in \text{prf}(D)$.

The second item is evident from the definition of the $\text{min-}S'$ -max-prf semantics, as presented in Definition 8. This definition states that the set of all $\text{min-}S'$ -max-prf for D is a subset of the set of all preferred interpretations of D .

The last two items are shown by our examples. \square

Returning to our induced ADFs, since every $\text{min-}S'$ -max preferred interpretation is admissible, we have the following result.

Corollary 5. Given KB $\mathcal{K} = \langle \mathcal{D}, \mathcal{P} \rangle$ and its induced ADF $D(\mathcal{K})$. $D(\mathcal{K})$ satisfies all three rational postulates for $\text{min-}S'$ -max-prf semantics (for any S').

4.3. Repair Semantics for KBs coincide with Minmax-Preferred Semantics for ADFs

Our final aim is to show that our induced ADFs can be used for inconsistency-tolerant query answering. In particular, we want to show that atomic query entailment under AR and brave semantics corresponds to skeptical and credulous acceptance under $\text{min-}S'$ -max preferred semantics for the case $S' = \mathcal{D}$, i.e. under $\text{min-}\mathcal{D}$ -max preferred semantics. For this, we need to show that $\text{min-}\mathcal{D}$ -max preferred interpretations do indeed capture the set of chases of ABox repairs.

We first show how admissible interpretations can be constructed from repairs.

Lemma 6. Given a KB $\mathcal{K} = \langle \mathcal{D}, \mathcal{P} \rangle$, $\mathcal{A}' \in \text{Repairs}(\mathcal{K})$, and let w be an interpretation of $D(\mathcal{K})$ s.t. for each $B \in S_{\mathcal{K}}$, $w(B) = \mathbf{t}$ if $B \in \text{chase}(\mathcal{A}')$, and $w(B) = \mathbf{f}$ otherwise. Then, w is admissible.

Proof. We have to show that $w(B) = \Gamma_{D(\mathcal{K})}(w)(B)$ if $w(B) \in \{\mathbf{t}, \mathbf{f}\}$. Let $\varphi_B = \varphi_B^1 \wedge \varphi_B^2 \wedge \varphi_B^3$ be the acceptance condition of B in $D(\mathcal{K})$, as introduced in Definition 6, where for convenience, we set $\varphi_B^2 = \top$ in case $B \in \mathcal{D}$.

Assume that $w(B) = \mathbf{t}$. To show that w is an admissible interpretation, we show that $\Gamma_{D(\mathcal{K})}(w)(B) = \mathbf{t}$. For this, we need to show that $\varphi_B^w \equiv \top$, in particular that $w(\varphi_B^1)^w \equiv (\varphi_B^2)^w \equiv (\varphi_B^3)^w \equiv \top$. Toward a contradiction, assume that $(\varphi_B)^w \not\equiv \top$. Since w is two-valued, then $(\varphi_B)^w \equiv \perp$, which means that $(\varphi_B^1)^w \equiv \perp$, $(\varphi_B^2)^w \equiv \perp$, or $(\varphi_B^3)^w \equiv \perp$.

- First, we assume that $(\varphi_B^1)^w \equiv \perp$. That is, there exists $B \wedge \bigwedge_{i=1}^n C_i \rightarrow D \in \mathcal{P}_{\mathcal{D}}^{\text{ch}}$ such that $(\bigwedge_{i=1}^n C_i \rightarrow D)^w \equiv \perp$, i.e., $w(C_i) = \mathbf{t}$ for each i , but $w(D) = \mathbf{f}$. However, if $w(C_i) = \mathbf{t}$, then by the definition of w , $C_i \in \text{chase}(\mathcal{A}')$. Since $C_i \in \text{chase}(\mathcal{A}')$ for each i , and $B \in \text{chase}(\mathcal{A}')$, our rule in \mathcal{P} applies and $D \in \text{chase}(\mathcal{A}')$. It follows that $w(D) = \mathbf{t}$ by construction. This is a contradiction to our assumption, so that $(\varphi_B^1)^w \equiv \perp$ cannot hold. Consequently, $(\varphi_B^1)^w \not\equiv \perp$.
- Next, assume that $(\varphi_B^2)^w \equiv \perp$. Then, for each $\bigwedge_{i=1}^n A_i \rightarrow B \in \mathcal{P}_{\mathcal{D}}^{\text{ch}}$, there must exist an i , s.t., $w(A_i) = \mathbf{f}$. By the definition of interpretation w , then $A_i \notin \text{chase}(\mathcal{A}')$. If for each $\bigwedge_{i=1}^n A_i \rightarrow B \in \mathcal{P}_{\mathcal{D}}^{\text{ch}}$, there exists i , s.t., $A_i \notin \text{chase}(\mathcal{A}')$, B cannot be derived, and $B \notin \text{chase}(\mathcal{A}')$, so that $w(B) = \mathbf{f}$. This is a contradiction by the assumption that $w(B) = \mathbf{t}$. We obtain that $(\varphi_B^2)^w \not\equiv \perp$.
- Finally, assume that $(\varphi_B^3)^w \equiv \perp$. That is, there exists $B \wedge E \rightarrow \perp \in \mathcal{P}_{\mathcal{D}}^{\text{ch}}$ s.t., $w(\neg E) = \mathbf{f}$, i.e., $w(E) = \mathbf{t}$. By the definition of interpretation w , $w(E) = \mathbf{t}$ implies that $E \in \text{chase}(\mathcal{A}')$. This is a contradiction by the assumption that \mathcal{A}' is a repair because $B \in \text{chase}(\mathcal{A}')$ and $E \in \text{chase}(\mathcal{A}')$ would then imply that \mathcal{A}' is inconsistent. We obtain $(\varphi_B^3)^w \not\equiv \perp$.

Thus, if $w(B) = \mathbf{t}$, then $(\varphi_B^1)^w \equiv (\varphi_B^2)^w \equiv (\varphi_B^3)^w \equiv \top$, which implies $\Gamma_{D(\mathcal{K})}(w)(B) = \mathbf{t}$.

Next we show that if $w(B) = \mathbf{f}$, then $\Gamma_{D(\mathcal{K})}(w)(B) = \mathbf{f}$. Assume $w(B) = \mathbf{f}$ and $\Gamma_{D(\mathcal{K})}(w)(B) = \mathbf{t}$. This means that

$(\varphi_B)^w \equiv \top$, and in particular, $(\varphi_B^2)^w \equiv \top$, which could be due to some rule $\bigwedge_{i=1}^n A_i \rightarrow B \in \mathcal{P}_D^{ch}$ s.t. $w(A_i) = \mathbf{t}$, or because $B \in \mathcal{D}$

First assume there is $\bigwedge_{i=1}^n A_i \rightarrow B \in \mathcal{P}_D^{ch}$ s.t. $w(A_i) = \mathbf{t}$. By our construction, $A_i \in \text{chase}(\mathcal{A}')$, which implies $B \in \text{chase}(\mathcal{A}')$ by definition of the chase. This contradicts $w(B) = \mathbf{f}$, and thus $(\varphi_B^2)^w$ must be contradictory, and consequently also $(\varphi_B)^w \equiv \perp$.

Assume otherwise that $B \in \mathcal{D}$. We show that this would imply that $\text{chase}(\mathcal{A}') \cup \{B\}$ is consistent, contrary to the assumption that \mathcal{A}' is a repair. We first show that $\text{chase}(\mathcal{A}') \cup \{B\} = \text{chase}(\mathcal{A}' \cup \{B\})$. Assume there is a rule $B \wedge C_1 \wedge \dots \wedge C_n \rightarrow D \in \mathcal{P}_D^{ch}$ that is applicable on $\text{chase}(\mathcal{A}') \cup \{B\}$, that is, $C_1, \dots, C_n \in \text{chase}(\mathcal{A}')$. From latter, it follows by construction that $w(C_i) = \mathbf{t}$ for all $1 \leq i \leq n$. Moreover, by construction of φ_B , $C_1 \wedge \dots \wedge C_n \rightarrow D$ occurs as conjunct in φ_B^1 , and since $(\varphi_B^3)^w \equiv \top$, we must have $w(D) = \mathbf{t}$ and $D \in \text{chase}(\mathcal{A}')$. That means, the conclusion of every rule that is applicable on $\text{chase}(\mathcal{A}') \cup \{B\}$ is already in $\text{chase}(\mathcal{A}')$, so that indeed $\text{chase}(\mathcal{A}' \cup \{B\}) = \text{chase}(\mathcal{A}') \cup \{B\}$. It remains to show that $\mathcal{A}' \cup \{B\}$ is consistent, for which we now only have to show that there is no $E \wedge B \rightarrow \perp \in \mathcal{P}_D^{ch}$ s.t. $E \in \text{chase}(\mathcal{A}')$. For any such integrity constraint, because $(\varphi_B^3)^w \equiv \top$ and φ_B^3 contains $\neg E$ as conjunct, $w(E) = \mathbf{f}$, and by construction of w , $E \notin \text{chase}(\mathcal{A}')$. We obtain that $\mathcal{A}' \cup \{B\}$ is consistent, contrary to the initial assumption that \mathcal{A}' is a repair. Correspondingly, the assumption that $w(B) = \mathbf{f}$ and $\Gamma_{D_{\mathcal{K}}}(w)(B) = \mathbf{t}$ cannot be true.

We obtain that $w = \Gamma_{D_{\mathcal{K}}}(w)$, and thus that it is an admissible interpretation of $D_{\mathcal{K}}$. \square

Using Lemma 6, we first show that, if we restrict our attention to those statements that correspond to database facts, then indeed do min- \mathcal{D} -max preferred interpretation capture the set of repairs. That is, the *maximizing* part of our semantics works as intended.

Lemma 7. *For every $\mathcal{A}' \in \text{Repairs}(\mathcal{K})$, there exists a min- \mathcal{D} -max preferred interpretation w of $D(\mathcal{K})$ s.t. $w^{\mathbf{t}}|_{\mathcal{D}} = \mathcal{A}'$, and for every min- \mathcal{D} -max preferred interpretation w of $D(\mathcal{K})$, there exists $\mathcal{A}' \in \text{Repairs}(\mathcal{K})$ s.t. $\mathcal{A}' = w^{\mathbf{t}}|_{\mathcal{D}}$.*

Proof. (\Rightarrow) Assume that $\mathcal{A}' \in \text{Repairs}(\mathcal{K})$. We show that $D(\mathcal{K})$ has a min- \mathcal{D} -max preferred interpretation w s.t., $\mathcal{A}' = w^{\mathbf{t}}|_{\mathcal{D}}$. We construct w based on \mathcal{A}' as in the proof for Lemma 6. The constructed interpretation is an admissible interpretation of $D(\mathcal{K})$. Furthermore, since each $B \in S_{\mathcal{K}}$ is either assigned to \mathbf{t} or \mathbf{f} , w is a preferred and two-valued model of $D(\mathcal{K})$.

It remains to show that w is a min- \mathcal{D} -max preferred interpretation of $D(\mathcal{K})$. Toward a contradiction, assume that w is not a min- \mathcal{D} -max preferred interpretation of $D(\mathcal{K})$. That is, either

1. $w^{\mathbf{t}}|_{\mathcal{D}}$ is not maximal among all preferred interpretations of $D(\mathcal{K})$, which means there exists $u \in \text{prf}(D(\mathcal{K}))$ s.t., $w^{\mathbf{t}}|_{\mathcal{D}} \subset u^{\mathbf{t}}|_{\mathcal{D}}$;
2. or, there exists $u \in \text{prf}$ s.t., $w^{\mathbf{t}}|_{\mathcal{D}} = u^{\mathbf{t}}|_{\mathcal{D}}$, but $u^{\mathbf{t}} \subset w^{\mathbf{t}}$.

We first show that Item 1 cannot occur. Toward a contradiction, we assume that there exists $u \in \text{prf}(D)$ s.t., $w^{\mathbf{t}}|_{\mathcal{D}} \subset u^{\mathbf{t}}|_{\mathcal{D}}$. Set $\mathcal{A}'' = u^{\mathbf{t}}|_{\mathcal{D}}$. By construction, $\mathcal{A}'' \subseteq \mathcal{D}$ and $\mathcal{A}' \subset \mathcal{A}''$. We can furthermore show that \mathcal{A}'' is

consistent: because u is preferred, by Theorem 1 we can apply Postulate 1 (closure), which by induction implies $\text{chase}(\mathcal{A}'') \subseteq u^{\mathbf{t}}$. Moreover, u must satisfy Postulate 2 (consistency), $\text{chase}(\mathcal{A}'')$ does not invalidate any integrity constraints. This means that \mathcal{A}'' is consistent, so that \mathcal{A}' cannot be maximal and consistent, and thus not a repair.

Next, we show that Item 2 cannot occur. Assume that there exist a preferred interpretation u s.t., $w^{\mathbf{t}}|_{\mathcal{D}} = u^{\mathbf{t}}|_{\mathcal{D}}$, but $u^{\mathbf{t}} \subset w^{\mathbf{t}}$. If $u^{\mathbf{t}} \subset w^{\mathbf{t}}$, then there exists $A \notin u^{\mathbf{t}}$, and $A \in w^{\mathbf{t}}$. By construction, $w^{\mathbf{t}} = \text{chase}(\mathcal{A})$. Furthermore, $\text{chase}(\mathcal{A}') = \text{chase}(w^{\mathbf{t}}|_{\mathcal{D}}) = \text{chase}(u^{\mathbf{t}}|_{\mathcal{D}})$, because $u^{\mathbf{t}}|_{\mathcal{D}} = w^{\mathbf{t}}|_{\mathcal{D}}$, by our assumption. We obtain that $A \in \text{chase}(u^{\mathbf{t}}|_{\mathcal{D}})$ but $A \notin u^{\mathbf{t}}$, so that $u^{\mathbf{t}} \subset \text{chase}(u^{\mathbf{t}}|_{\mathcal{D}})$. This means there must be rules in \mathcal{P} that are applicable on $u^{\mathbf{t}}$ but not applied, which contradicts Postulate 1, which by Theorem 1 means that u cannot be a preferred extension, contrary to our initial assumption that $u \in \text{prf}(D)$.

(\Leftarrow) Assume that w is a min- \mathcal{D} -max preferred interpretation. Let $\mathcal{A}' = w^{\mathbf{t}}|_{\mathcal{D}}$. We show that $\mathcal{A}' \in \text{Repairs}(\mathcal{K})$, for which we need to show that \mathcal{A}' is consistent and a maximal subset of \mathcal{D} with that property.

First, we show that \mathcal{A}' is consistent. If \mathcal{A}' were inconsistent, then there would be an integrity constraint $A \wedge B \rightarrow \perp \in \mathcal{P}_D^{ch}$ s.t. $A, B \in \text{chase}(\mathcal{A}')$. By repeated application of Postulate 1 (closure), we obtain that $\text{chase}(\mathcal{A}') = \text{chase}(w^{\mathbf{t}}|_{\mathcal{D}}) \subseteq w^{\mathbf{t}}$, which would imply that $A, B \in w^{\mathbf{t}}$, contradicting Postulate 2 (consistency).

Next we have to show that \mathcal{A}' is maximal. Toward a contradiction assume that \mathcal{A}' is not maximal. Thus, there exists \mathcal{A}'' s.t., $\mathcal{A}' \subset \mathcal{A}''$ and \mathcal{A}'' is consistent.

In the same way as in the first part of this proof, we can construct a min- \mathcal{D} -max preferred interpretation u s.t. $u|_{\mathcal{D}} = \mathcal{A}''$. But then, we have found an interpretation $u \in \text{prf}(D(\mathcal{K}))$ s.t. $w^{\mathbf{t}}|_{\mathcal{D}} \subset u^{\mathbf{t}}|_{\mathcal{D}}$, contradicting that w is min- \mathcal{D} -max preferred interpretation. \square

To be able to answer queries, which rely on derived facts, we need to also make sure that the chase is correctly reflected in our interpretations. Intuitively, we have to show that also the *minimizing* part of our interpretations works as intended.

Lemma 8. *For every $\mathcal{A}' \in \text{Repairs}(\mathcal{K})$, there exists a min- \mathcal{D} -max preferred interpretation w of $D(\mathcal{K})$ s.t. $\text{chase}(\mathcal{A}') = w^{\mathbf{t}}$, and for every min- \mathcal{D} -max preferred interpretation w , there exists $\mathcal{A}' \in \text{Repairs}(\mathcal{K})$ s.t. $\text{chase}(\mathcal{A}') = w^{\mathbf{t}}$.*

Proof. The first direction follows directly from constructions we already used: the interpretation constructed in Lemma 6 based on any repair \mathcal{A}' satisfies $\text{chase}(\mathcal{A}') = w^{\mathbf{t}}$, and is shown to be min- \mathcal{D} -max preferred interpretation.

For the other direction, let w be a min- \mathcal{D} -max preferred interpretation of $D(\mathcal{K})$. By Lemma 7 there exists $\mathcal{A}' \in \text{Repairs}(\mathcal{K})$ s.t. $\mathcal{A}' = w^{\mathbf{t}}|_{\mathcal{D}}$. Applying Postulate 1 (closure), we obtain $\text{chase}(\mathcal{A}') \subseteq w^{\mathbf{t}}$. By the previous direction, there exists a min- \mathcal{D} -max preferred interpretation v s.t., $\text{chase}(\mathcal{A}') = v^{\mathbf{t}}$. We obtain $v^{\mathbf{t}} \subseteq w^{\mathbf{t}}$, which by the minimality condition of the Definition 8 (Item 2), implies $w^{\mathbf{t}} = v^{\mathbf{t}}$. Hence, $\text{chase}(\mathcal{A}') = w^{\mathbf{t}}$. \square

Corollary 9. *For every consistent KB \mathcal{K} , the induced ADF has a unique min- \mathcal{D} -max preferred interpretation.*

Proof. This follows from the fact a consistent KB has exactly one repair, so that every min- \mathcal{D} -max preferred semantics would have to capture exactly its chase. \square

Note that the converse of Corollary 9 does not hold: consider the KB $\mathcal{K} = \langle \mathcal{D}, \mathcal{P} \rangle$ s.t., $\mathcal{D} = \{A(a), B(a)\}$ and $\mathcal{P} = \{A(x) \rightarrow \perp\}$. \mathcal{K} has one repair, the induced ADF exactly one min- \mathcal{D} -max preferred interpretation, but \mathcal{K} is inconsistent.

Theorem 10. *Given a KB $\mathcal{K} = \langle \mathcal{D}, \mathcal{P} \rangle$ a query $q(\vec{x})$ and a vector \vec{c} of constants the length of \vec{x} . Let $D(\mathcal{K})$ be the induced ADF of \mathcal{K} .*

1. \vec{c} is an answer of $q(\vec{x})$ under AR semantics, iff $Skept_{min-\mathcal{D}-max-prf}(q(\vec{c}), D(\mathcal{K}))$ is satisfied.
2. \vec{c} is an answer of $q(\vec{x})$ under brave semantics, iff $Cred_{min-\mathcal{D}-max-prf}(q(\vec{c}), D(\mathcal{K}))$ is satisfied.

Proof. By Theorem 8, for every $\mathcal{A}' \in \text{Repairs}(\mathcal{K})$, there exists a min- \mathcal{D} -max preferred interpretation v of $D(\mathcal{K})$ s.t. $\text{chase}(\mathcal{A}') = v^{\mathbf{t}}$. (and vice versa). Therefore, \vec{c} is an answer of $q(\vec{x})$ under AR semantics, i.e., for every repair $\mathcal{A}' \in \text{Repairs}(\mathcal{K})$, $\langle \mathcal{P}, \mathcal{A}' \rangle \models q(\vec{c})$, iff for each min- \mathcal{D} -max preferred interpretation v (of $D(\mathcal{K})$), $v(q(\vec{c})) = \mathbf{t}$, since $\text{chase}(\mathcal{A}') = v^{\mathbf{t}}$, i.e., iff $Skept_{min-\mathcal{D}-max-prf}(q(\vec{c}), D(\mathcal{K}))$ is satisfied.

Similarly, \vec{c} is considered an answer of $q(\vec{x})$ under brave semantics if for some repair $\mathcal{A}' \in \text{Repairs}(\mathcal{K})$, $\langle \mathcal{P}, \mathcal{A}' \rangle \models q(\vec{c})$, if and only if there exists an min- \mathcal{D} -max preferred interpretation v (of $D(\mathcal{K})$), such that $v(q(\vec{c})) = \mathbf{t}$, since $\text{chase}(\mathcal{A}') = v^{\mathbf{t}}$, i.e., iff $Cred_{min-\mathcal{D}-max-prf}(q(\vec{c}), \mathbf{t}, D(\mathcal{K}))$ is satisfied. \square

5. Complexity

We first analyze the complexity of our new semantics for the standard reasoning tasks of ADFs [33], before looking at the specific ADFs that are induced from KBs through our construction in Section 5.2.

We make use of standard complexity classes of the polynomial hierarchy [34]. In particular, P is the class of decision problems decidable in polynomial time, NP is the class of decision problems decidable in polynomial time with a non-deterministic algorithm. The class Σ_i^P contains all decision problems decidable in non-deterministic polynomial time with access to an oracle in Σ_{i-1}^P (i.e., a problem in Σ_{i-1}^P can be decided in constant time), for $i > 0$, $\Sigma_0^P = P$ and $\Sigma_1^P = NP$. Class Π_i^P is the complementary class of Σ_i^P .

5.1. Complexity of min- S' -max preferred semantics on ADFs

Theorem 11. *It holds that $Ver_{min-S'-max-prf}$ is in Π_3^P for ADFs.*

Proof. Consider an arbitrary instance of the problem, that is, a given ADF $D = (S, L, C)$, a three-valued interpretation v and $S' \subseteq S$. We show membership in Σ_3^P for the complementary problem of v not being a min- S' -max preferred interpretation.

The algorithm guesses an interpretation w , checks $w \in prf(D)$ (in Π_2^P [32]), and whether any of the following is satisfied:

1. $v \leq_i w$
2. $v^{\mathbf{t}}|_{S'} \subset w^{\mathbf{t}}|_{S'}$
3. $w|_{S'} = v|_{S'}$ and $w^{\mathbf{t}} \subset v^{\mathbf{t}}$

σ	$Cred_{\sigma}$	$Skept_{\sigma}$	Ver_{σ}
<i>adm</i>	$\Sigma_2^P\text{-c}$	trivial	coNP-c
<i>prf</i>	$\Sigma_2^P\text{-c}$	$\Pi_3^P\text{-c}$	$\Pi_2^P\text{-c}$
<i>min-S'-max-prf</i> (for induced ADFs)	in Σ_4^P in coNP	in Π_4^P in NP	$\Pi_3^P\text{-c}$ in P

Table 2

Complexity under min- S' -max preferred semantics of ADFs and of the induced ADF from a KB.

If any of the above items is successful, then $v \notin min-S'-max-prf(D)$. Thus, the complementary problem $Ver_{min-S'-max-prf}(v, D)$ is in Σ_3^P . From this it follows that the verification problem is in Π_3^P for ADFs. \square

Theorem 12. *It holds that $Cred_{min-S'-max-prf}$ is in Σ_4^P for ADFs.*

Proof. Let $D = (S, L, C)$ be an arbitrary ADF, $s \in S$. To check if there exists a min- S' -max preferred interpretation v satisfying $v(s) = \mathbf{t}$, guess an interpretation v with $v(s) = \mathbf{t}$, and then verify whether $v \in min-S'-max-prf(D)$. According to Theorem 11, verifying whether $v \in min-S'-max-prf(D)$ is in Π_3^P . Thus, the combined guessing and checking process results in $NP^{\Pi_3^P} = \Sigma_4^P$. \square

Theorem 13. *It holds that $Skept_{min-S'-max-prf}$ is in Π_4^P for ADFs.*

Proof. Given an ADF $D = (S, R, C)$, and a statement s , answering $Skept_{min-S'-max-prf}(s, D)$ involves considering the complementary problem. In this case, we determine whether there exists a min- S' -max preferred interpretation v in which s is not assigned to \mathbf{t} . As per Theorem 11, the task of checking if v is a min- S' -max preferred interpretation of D is in Π_3^P . Thus, $Skept_{min-S'-max-prf}(s, D)$ is in Π_4^P . \square

5.2. Complexity of min- \mathcal{D} -max preferred semantics on the induced ADFs

The hardness results in the previous subsection all depend on the fact that already the admissibility-problem is coNP-complete for ADFs. The ADFs we define in Definition 6 have a specific syntactic form that allows us to show admissibility in polynomial time.

Lemma 14. *The verification problem under admissible semantics of ADFs $D = (S, L, C)$ induced by KBs is in P.*

Proof. To decide whether v is admissible, we have to determine whether $v \leq_i \Gamma_D(v)$. We show that $\Gamma_D(v)$ can be computed in polynomial time, which shows that this test is computed in polynomial time as well.

To compute $\Gamma_D(v)$, we need to determine, for each $A \in S$, whether $v(\varphi_A)$ is irrefutable, unsatisfiable, or contingent (Definition 2). Fix some such A . We describe a deterministic procedure to compute $\Gamma_D(v)(A)$. Roughly, we try to find a disjunction φ_A^2 which is irrefutable or contingent together with $\varphi_A^1 \wedge \varphi_A^2$. We describe the procedure for each disjunct in the following. Each such disjunct is of the form $\bigwedge_{i=1}^n B_i$ and corresponds to a rule $\bigwedge_{i=1}^n B_i \rightarrow A \in \mathcal{P}_D^{ch}$. To check satisfiability, we attempt to extend v to a satisfying interpretation w . If for some $i \in \{1, \dots, n\}$, $v(B_i) = \mathbf{f}$, we know that this disjunct cannot be satisfied in any extension

of v . Otherwise, we build a partial valuation v_0 based on v by setting $v_0(B_i) = \mathbf{t}$ for every $1 \leq i \leq n$, $v_0(A') = v$ for every other atom A' s.t. $v(A') \neq \mathbf{u}$, and $v(A') = \mathbf{u}$ for the remaining atoms. Then we iteratively go through the rules in $\mathcal{P}_D^{\text{ch}}$ that are relevant for the acceptance condition. In particular, we try to generate a sequence of partial valuations v_0, \dots, v_m as follows, where $i \geq 0$:

1. If φ_A^1 contains an implication $A_1 \wedge \dots \wedge A_n \rightarrow B$ s.t. $v_i(A_i) = \mathbf{t}$ for $1 \leq i \leq n$ and $v_i(B) = \mathbf{u}$, continue with v_{i+1} which extends v_i by $v_{i+1}(B) = \mathbf{t}$.
2. If φ_A^3 contains a conjunct $\neg E$ with $v_i(E) = \mathbf{t}$, we know that the current disjunct cannot be satisfied.
3. If φ_A^1 contains an implication $A_1 \wedge \dots \wedge A_n \rightarrow B$ s.t. $v_i(A_i) = \mathbf{t}$ for $1 \leq i \leq n$ and $v_i(B) = \mathbf{f}$, we know that the current disjunct cannot be satisfied.

If for none of the disjuncts, we are successful, we know that φ_A^v cannot be satisfied by any valuation, and we can set $\Gamma_D(v)(A) = \mathbf{f}$. If we manage to find a partial valuation v_m on which no more steps can be applied, we found a partial way of satisfying the acceptance condition, which means we know that φ_A^v is satisfiable, but not yet whether it is also irrefutable. For this, we just need to check whether there is some atom A' s.t. $v(A') \neq v_m(A')$. If there is not, our process confirmed that all conjuncts of φ_A^v are already satisfied by v , and we can set $\Gamma_D(v)(A) = \mathbf{t}$. If there is, we can find a valuation v' that makes $\varphi_A^1 \wedge \varphi_A^3$ with the current disjunct false by setting $v'(A) = \mathbf{f}$ and $v'(A') = v_m(A')$ for all other atoms. If we cannot find a disjunct that justifies $v'(A) = \mathbf{t}$, but showed satisfiability, we consequently can set $\Gamma_D(v)(A) = \mathbf{u}$. \square

But indeed, we know that more complexities go down:

Theorem 15. *For ADFs induced by KBs, $Ver_{\min\text{-}D\text{-max}\text{-prf}}$ is in Π_2^P , $Cred_{\min\text{-}D\text{-max}\text{-prf}}$ in NP and $Skept_{\min\text{-}D\text{-max}\text{-prf}}$ in coNP.*

Proof. Consider an arbitrary instance of the problem, that is, a KB $\mathcal{K} = \langle \mathcal{D}, \mathcal{P} \rangle$, the induced ADF $D(\mathcal{K})$, and a three-valued interpretation v (for D). By Lemma 8, if v is a min- \mathcal{D} -max preferred interpretation, then there exists $\mathcal{A}' \in \text{Repairs}(\mathcal{K})$ s.t. $\text{chase}(\mathcal{A}') = v^{\mathbf{t}}$. Instead of checking if v is a min- \mathcal{D} -max preferred interpretation, we check if there is a repair \mathcal{A}' s.t. $\text{chase}(\mathcal{A}') = v^{\mathbf{t}}$. In this case, by Lemma 7, $\mathcal{A}' = v^{\mathbf{t}}|_{\mathcal{D}}$ is a repair. Thus, first we check if $v^{\mathbf{t}}|_{\mathcal{D}}$ is a repair. If it is, we check if $\text{chase}(v^{\mathbf{t}}|_{\mathcal{D}}) = v^{\mathbf{t}}$. Both steps can be done in polynomial time. Therefore, $Ver_{\min\text{-}D\text{-max}\text{-prf}}$ is in P.

For $Cred_{\min\text{-}D\text{-max}\text{-prf}}$, to check if there exists min- \mathcal{D} -max preferred interpretation v satisfying $v(s) = \mathbf{t}$, guess an interpretation with $v(s) = \mathbf{t}$ and verify if $v \in \text{min-}\mathcal{D}\text{-max-prf}(D)$. Verification is in P. Therefore, $Cred_{\min\text{-}D\text{-max}\text{-prf}}$ is in NP. For $Skept_{\min\text{-}D\text{-max}\text{-prf}}$ we consider the complementary problem, i.e., we check if there is a $v \in \text{min-}\mathcal{D}\text{-max-prf}(D)$, s.t. $v(s) \neq \mathbf{t}$. \square

While for $Ver_{\min\text{-}S'\text{-max}\text{-prf}}$ we only have an coNPUpper bound, we observe that the problem becomes polynomial if we focus on the positive part only: given v , we can decide whether $v^{\mathbf{t}} = w^{\mathbf{t}}$ for some $w \in \text{min-}\mathcal{D}\text{-max-prf}$ as follows: determine whether $v^{\mathbf{t}}|_{\mathcal{D}}$ is a repair, verify $v^{\mathbf{t}} = \text{chase}(v^{\mathbf{t}}|_{\mathcal{D}})$ (using Lemma 8).

We can use Theorem 15 directly to (re-)prove upper bounds for various settings of inconsistency-tolerant reasoning. As is common, we look at the decision problem

corresponding to query answering, which is the *query entailment problem* which checks whether for a given instantiated query $q(\vec{c})$, $\mathcal{K} \models q(\vec{c})$. Recall that for *data complexity*, we assume the size of the program is fixed, while for *combined complexity* we assume both database and program to be given by the input.

Theorem 16. *Atomic query entailment under AR/brave semantics is in coNP/NP*

- *data complexity for datalog programs and acyclic Horn-SRI programs,*
- *combined complexity for datalog if the arity of predicates is bounded.*

Proof. The first results rely only on the size of $\text{chase}(\mathcal{A})$, since the size of the induced ADF is polynomial in it. For datalog with bounded arity of predicates (e.g. bounded by k), we observe that $\text{chase}(\mathcal{A})$ is polynomial in \mathcal{A} since every term in $\text{chase}(\mathcal{A})$ also occurs in \mathcal{A} , and the number of atoms one can build from l terms is $m \times l^k$, with m the number of predicates that occur in total. The same statement gives us tractability in data complexity for datalog in general, since with a fixed program, the arity is trivially bounded as well. For Horn-SRI, we observe by inspection of the rules in Table 1 that for a fixed acyclic program, each Skolem term contains exactly one variable, and consequently, each Skolem term in $\text{chase}(\mathcal{A})$ contains exactly one constant, and its nesting depth is determined only by \mathcal{P} . Consequently, the number of terms in $\text{chase}(\mathcal{A})$, as well as of atoms, is also polynomial in the size of \mathcal{A} . Note also that in all cases, the construction of the chase runs in polynomial time as well by simple forward-chaining. \square

6. Conclusion

We proposed a method for using ADFs to answer queries over inconsistent KBs. ADFs are useful for this task because they offer an abstract way to represent statements and arguments. To construct our framework, we consider each atom in the chase of a KB as an abstract node or statement. This eliminates the need to explicitly build statements within the ADF. We showed that the induced ADFs satisfy rationality postulates, and introduced a new semantics for ADFs that is needed to reproduce inconsistency-tolerant semantics for KBs. Finally, we gave some complexity upper bounds for this new semantics, which in the general case are higher than for preferred semantics, but substantially lower for the ADFs obtained with our construction.

In the future, we want to improve the complexity bounds and provide matching lower bounds for the general case of our new semantics. Furthermore, we want to investigate how our ADFs can be used to provide explanations of query answers for inconsistency-tolerant KBs to end-users. An idea could be to extend the graphical, interactive proof-exploration tool EVONNE [1], to explore the structure of the ADF induced by the KB starting from the atom corresponding to the query answer.

Acknowledgments

This publication has been supported by the Netherlands Organisation for Scientific Research (NWO) through the Hybrid Intelligence Gravitation Programme with project number 024.004.022

References

- [1] J. Méndez, C. Alrabbaa, P. Koopmann, R. Langner, F. Baader, R. Dachselt, Evonne: A visual tool for explaining reasoning with OWL ontologies and supporting interactive debugging, *Comput. Graph. Forum* 42 (2023).
- [2] C. Alrabbaa, S. Borgwardt, P. Koopmann, A. Kovtunova, Explaining ontology-mediated query answers using proofs over universal models, in: *International Joint Conference on Rules and Reasoning*, Springer, 2022, pp. 167–182.
- [3] A. Elhalawati, M. Krötzsch, S. Mennicke, An existential rule framework for computing why-provenance on-demand for datalog, in: *International Joint Conference on Rules and Reasoning*, 2022, pp. 146–163.
- [4] A. Cali, G. Gottlob, T. Lukasiewicz, A general datalog-based framework for tractable query answering over ontologies, *J. Web Semant.* 14 (2012) 57–83.
- [5] M. Ortiz, S. Rudolph, M. Simkus, Query answering in the Horn fragments of the description logics *SHOIQ* and *SRQIQ*, in: *Proceedings of IJCAI 2011*, 2011, pp. 1039–1044.
- [6] A. Borgida, D. Calvanese, M. Rodriguez-Muro, Explanation in the DL-Lite family of description logics, in: *On the Move to Meaningful Internet Systems: OTM 2008 Confederated International Conferences*, Springer, 2008, pp. 1440–1457.
- [7] T. Lukasiewicz, E. Malizia, C. Molinaro, Explanations for inconsistency-tolerant query answering under existential rules, in: *AAAI*, 2020, pp. 2909–2916.
- [8] C. Alrabbaa, F. Baader, S. Borgwardt, P. Koopmann, A. Kovtunova, Finding small proofs for description logic entailments: Theory and practice, in: *LPAR 2020*, 2020, pp. 32–67.
- [9] M. Bienvenu, C. Bourgaux, Querying and repairing inconsistent prioritized knowledge bases: Complexity analysis and links with abstract argumentation, in: *Proceedings of KR*, 2020, pp. 141–151.
- [10] G. Brewka, S. Ellmauthaler, H. Strass, J. P. Wallner, S. Woltran, Abstract dialectical frameworks. An overview, *IFCoLog Journal of Logics and their Applications (FLAP)* 4 (2017).
- [11] M. Bienvenu, C. Bourgaux, F. Goasdoué, Computing and explaining query answers over inconsistent DL-Lite knowledge bases, *J. Artif. Intell. Res.* 64 (2019) 563–644.
- [12] A. Arioua, N. Tamani, M. Croitoru, Query answering explanation in inconsistent datalog +/- knowledge bases, in: *Proceedings of DEXA 2015*, Springer, 2015, pp. 203–219.
- [13] H. Strass, Instantiating knowledge bases in abstract dialectical frameworks, in: *Proceedings of CLIMA*, Springer, 2013, pp. 86–101.
- [14] A. Z. Wyner, T. J. M. Bench-Capon, P. E. Dunne, On the instantiation of knowledge bases in abstract argumentation frameworks, in: *Proceedings of CLIMA*, Springer, 2013, pp. 34–50.
- [15] S. Abiteboul, R. Hull, V. Vianu, *Foundations of databases*, volume 8, Addison-Wesley Reading, 1995.
- [16] D. Calvanese, G. De Giacomo, Data integration: A logic-based perspective, *AI magazine* 26 (2005) 59–59.
- [17] A. K. Chandra, P. M. Merlin, Optimal implementation of conjunctive queries in relational data bases, in: *Proceedings of the ninth annual ACM symposium on Theory of computing*, 1977, pp. 77–90.
- [18] M. Caminada, Rationality postulates: Applying argumentation theory for non-monotonic reasoning, *FLAP* 4 (2017).
- [19] O. Arieli, A. Borg, C. Straßer, A postulate-driven study of logical argumentation, *Artif. Intell.* 322 (2023) 103966.
- [20] P. M. Dung, On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games, *Artif. Intell.* 77 (1995) 321–358.
- [21] P. Baroni, M. Caminada, M. Giacomin, An introduction to argumentation semantics, *Knowledge Engineering Review* 26 (2011) 365–410.
- [22] O. Cocarascu, A. Rago, F. Toni, Extracting dialogical explanations for review aggregations with argumentative dialogical agents, in: *AAMAS, International Foundation for Autonomous Agents and Multiagent Systems*, 2019, pp. 1261–1269.
- [23] A. J. García, C. I. Chesñevar, N. D. Rotstein, G. R. Simari, Formalizing dialectical explanation support for argument-based reasoning in knowledge-based systems, *Expert Syst. Appl.* 40 (2013) 3233–3247.
- [24] M. Caminada, L. Amgoud, On the evaluation of argumentation formalisms, *Artif. Intell.* 171 (2007) 286–310.
- [25] M. Caminada, Y. Wu, On the limitations of abstract argumentation, in: *Proceedings of the 23rd Benelux Conference on Artificial Intelligence (BNAIC 2011)*, 2011, pp. 59–66.
- [26] O. Arieli, A. Borg, C. Straßer, Tuning logical argumentation frameworks: A postulate-derived approach, in: *FLAIRS, AAAI Press*, 2020, pp. 557–562.
- [27] O. Arieli, A. Borg, C. Straßer, Characterizations and classifications of argumentative entailments, in: *KR*, 2021, pp. 52–62.
- [28] G. Brewka, S. Ellmauthaler, H. Strass, J. P. Wallner, S. Woltran, Abstract dialectical frameworks: An overview, in: *Handbook of Formal Argumentation*, College Publications, London, 2018, pp. 237–285.
- [29] G. Brewka, S. Woltran, Abstract dialectical frameworks, in: *Proc. KR*, 2010, pp. 102–111.
- [30] B. C. Grau, I. Horrocks, M. Krötzsch, C. Kupke, D. Magka, B. Motik, Z. Wang, Acyclicity notions for existential rules and their application to query answering in ontologies, *Journal of Artificial Intelligence Research* 47 (2013) 741–808.
- [31] T. Eiter, T. Lukasiewicz, L. Predoiu, Generalized consistent query answering under existential rules, in: *KR, AAAI Press*, 2016, pp. 359–368.
- [32] H. Strass, J. P. Wallner, Analyzing the computational complexity of abstract dialectical frameworks via approximation fixpoint theory, *Artif. Intell.* 226 (2015) 34–74.
- [33] W. Dvořák, P. E. Dunne, Computational problems in formal argumentation and their complexity, in: *Handbook of Formal Argumentation*, College Publications, London, 2018, pp. 631–687.
- [34] L. J. Stockmeyer, The polynomial-time hierarchy, *Theor. Comput. Sci.* 3 (1976) 1–22.