

Extending Abstract Argumentation Frameworks with Knowledge Bases

Gianvincenzo Alfano , Sergio Greco , Cristian Molinaro , Francesco Parisi , Irina Trubitsyna

Department of Informatics, Modeling, Electronics and System Engineering,
University of Calabria, Italy

{g.alfano, greco, c.molinaro, fparisi, trubitsyna}@dimes.unical.it

Abstract

Dung’s abstract Argumentation Framework (AF) has been extended in several directions to make knowledge representation and reasoning more intuitive and expressive. In this paper, we present the *Knowledge-based Argumentation Framework* (KAF), an extension of AF with a Knowledge Base (KB) expressed in *DL-Lite*, which includes concept and role instances describing the topology of an AF, besides additional knowledge on the domain. The KAF semantics is given by a set of KAF *extensions*, each consisting of an extension of the underlying AF together with a “pertinent” subset of the original KB, which is obtained by discarding assertions referring to arguments that have been ruled out in the AF extension. Then, the framework is further expanded into the *Constrained KAF* (CKAF), where a set of restricted relational calculus formulae is used for reasoning over ‘feasible’ subframeworks that satisfy the formulae and *minimally* differ from the original framework. We thoroughly investigate the computational complexity of classical reasoning problems under popular argumentation semantics, and show that well-known AF-based frameworks are special cases of CKAF.

1 Introduction

Formal argumentation has become a prominent research field in the area of knowledge representation and reasoning (Gabbay et al. 2021), with potential applications in several contexts, including, e.g., modeling dialogues, negotiation (Amgoud, Dimopoulos, and Moraitis 2007; Dimopoulos, Maily, and Moraitis 2019), and persuasion (Prakken 2009), thus promoting promising cross-field connections (Gabbay et al. 2024). At the heart of this field lies Dung’s *abstract Argumentation Framework* (AF), a foundational yet expressive formalism for representing disputes among agents (Dung 1995). An AF consists of a set of abstract arguments and a binary attack relation that defines how arguments interact: intuitively, if argument a attacks argument b , then b is considered acceptable only if a is not. Arguments themselves are treated as abstract entities whose acceptability essentially depends on the attack relation. This framework can be naturally represented as a directed graph, where nodes represent arguments and edges denote attacks.

The formal meaning of an AF is given in terms of *argumentation semantics*, which intuitively tell us the sets of arguments (called *extensions*) that can collectively be accepted. Several (argumentation) semantics have been in-




Figure 1: AF Λ of Example 1 (highlighted in blue), and (representation of the) assertions of the (C)KAF of our running example.

troduced, such as *grounded* (gr), *complete* (co), *preferred* (pr), and *stable* (st) (Dung 1995), each defining a set of σ -extensions, with $\sigma \in \{\text{gr}, \text{co}, \text{pr}, \text{st}\}$.

To enhance the expressive power of Dung’s framework, various extensions of the original framework have been proposed. These extensions include: Bipolar argumentation frameworks, adding *supports* alongside attacks (Nouioua and Risch 2011; Nouioua 2013; Villata et al. 2012); Frameworks with *recursive* attacks and supports (Cohen et al. 2015; Gottifredi et al. 2018; Cayrol et al. 2018; Alfano et al. 2024b); Dialectical frameworks (Brewka and Woltran 2010; Brewka et al. 2013); AFs with preferences (Amgoud and Cayrol 1998; Modgil and Prakken 2013; Alfano et al. 2023) and constraints (Coste-Marquis, Devred, and Marquis 2006; Arieli 2015; Alfano et al. 2024a); Claim-based AF (Dvorák, Rapberger, and Woltran 2020; Dvorák and Woltran 2020; Dvorák et al. 2023), where *claims* are associated to arguments and become the objects of reasoning.

Another important extension of AF concerns allowing the possibility of representing uncertainty about the existence of arguments and attacks. This has been carried out by either considering quantified uncertainty, thus combining formal argumentation with probability theory in the Probabilistic Argumentation Framework (PrAF) (Dung and Thang 2010; Li, Oren, and Norman 2011; Hunter 2012), or considering unquantified uncertainty by explicitly denoting the elements (arguments and attacks) which are uncertain (Baumeister et al. 2018; Baumeister et al. 2021), leading to the Incomplete Argumentation Framework (iAF).

Although significant advancements have been made in extending AF, there still remain situations where supplementary but potentially valuable information is neglected.

Example 1. Assume that there is a debate with

several participants expressing opinions about climate change. Opinions and their relationships are represented by means of AF $\Lambda_1 = \langle \{a, b, c, d, e\}, \{(a, b), (b, a), (c, b), (c, d), (d, c), (e, c)\} \rangle$ shown in Figure 1. We also know that the opinions synthesized by arguments a, b, c , and d are expressed by *john*, *mary*, *frank*, and *anne*, respectively; we do not know who expressed opinion e . AF Λ_1 has three complete extensions: $E_1 = \{d, e\}$ (which is also grounded), and $E_2 = \{a, d, e\}$ and $E_3 = \{b, d, e\}$ (which are also stable and preferred). $\square$

By modeling a problem through a plain AF, we may neglect important information that can be very useful. In our example, AF Λ_1 disregards useful information about the text associated with each argument, the author (and her gender and age), etc. Generally, a concrete argument is associated with a sentence with a proper meaning and with several additional information that could be very important in answering queries and making informed decisions.

This means that, in adopting the abstraction of AF, a large amount of relevant information may be lost, such as the argument's text, the author's id and gender, the date it has been introduced, the topic (e.g., what the argument is about), the polarity (e.g., whether the sentence conveys a positive, negative or neutral sentiment), the polarity rating, the sentence style (e.g., veracity, sarcasm, irony), the city and nationality of the author, and much more. An argumentation system should also carry auxiliary yet important information that could be very useful for users. That is, an argumentation framework should be coupled with a Knowledge Base (KB) containing information related to arguments. Prominent (families of) languages for such kind of KBs are those in Description Logics (DLs) (Brachman and Levesque 2004), where a KB is a pair consisting of the ABox component, containing the data, and the TBox component, containing axioms (enabling further information to be derived).

Example 2. Continuing with Example 1, assume now to have a KB containing in the ABox information on the topology of the AF (i.e., $\text{arg}(a)$, $\text{att}(a, b)$, and so on) as well as additional information such as authorship (i.e., $\text{author}(a, \text{john})$, $\text{author}(b, \text{mary})$, $\text{author}(c, \text{frank})$, and $\text{author}(d, \text{anne})$), the city of authors (i.e., $\text{resident}(\text{john}, \text{paris})$, $\text{resident}(\text{mary}, \text{paris})$, $\text{resident}(\text{frank}, \text{bath})$, and $\text{resident}(\text{anne}, \text{bath})$), and the nation cities are in (e.g., $\text{nation}(\text{paris}, \text{france})$ and $\text{nation}(\text{bath}, \text{uk})$).

Assume that the TBox contains two axioms (expressed according to the *DL-Lite* syntax) of the form: *i*) $\text{arg} \sqsubseteq \exists \text{author}$ (equivalent to the first-order logic formula $\forall x. \text{arg}(x) \Rightarrow \exists y. \text{author}(x, y)$), stating that every argument must have an author, and *ii*) $\exists \text{author}^- \sqsubseteq \exists \text{resident}$ (equivalent to the first-order logic formula $\forall x, y. \text{author}(x, y) \Rightarrow \exists z. \text{resident}(y, z)$), stating that every author must be resident in some place.

Using such a *DL-Lite* KB, the (certain) answer of the query $\text{author}(x, \text{john})$, asking for the arguments authored by *john*, gives as output argument a . Because of the axioms in the TBox, the answer of the query $\exists y, z. \text{author}(e, y), \text{resident}(y, z)$ is true, even if the

ABox does not contain the author of e . $\square$

Example 2 shows a possible integration of AF and KBs. However, as argumentation semantics are not leveraged during the reasoning process, the information on the AF's topology is not adequately exploited, although it is encoded into the ABox component. Thus, the classical reasoning loses the potential of AF in managing conflicts and uncertainty.

In this paper, we present a novel framework called *Knowledge-based Argumentation Framework* (KAF), augmenting AF with a *DL-Lite* KB (Calvanese et al. 2007), where the semantics is given by a set of KAF σ -extensions, each consisting of a σ -extension of the underlying AF together with a “pertinent” subset of the original KB. The latter is obtained by discarding assertions referring to arguments that have been ruled out in the AF extension. Notably, a KAF σ -extension (with $\sigma \in \{\text{gr}, \text{co}, \text{pr}, \text{st}\}$) can be represented as a *DL-Lite* KB.

Example 3. Continuing again with our running example, the (grounded) KAF extension yielded by $E_1 = \{d, e\}$ consists of the ABox containing $\text{arg}(d)$, $\text{arg}(e)$, $\text{author}(d, \text{anne})$, $\text{resident}(\text{john}, \text{paris})$, $\text{resident}(\text{mary}, \text{paris})$, $\text{resident}(\text{frank}, \text{bath})$, $\text{resident}(\text{anne}, \text{bath})$, $\text{nation}(\text{paris}, \text{france})$, and $\text{nation}(\text{bath}, \text{uk})$, that is the original ABox except assertions referring to arguments which are not in the extension E_1 , while the TBox is the one of the original KB. $\square$

Similarly to classical reasoning in AF, credulous and skeptical acceptance in KAF are defined over the set of KAF σ -extensions, consisting of *DL-Lite* KBs, each of them queried according to the usual notion of certain answer (Calvanese et al. 2007). For instance, the query $\exists x, y. \text{author}(x, y) \wedge \text{resident}(y, \text{bath})$, asking whether there exists an author of an argument resident in *bath*, is true under both credulous and skeptical reasoning, w.r.t. all the aforementioned argumentation semantics (as argument d , whose author is *anne* who is from *bath*, is a certain answer in every KAF σ -extension, with $\sigma \in \{\text{gr}, \text{co}, \text{pr}, \text{st}\}$, and the set of σ -extensions is non-empty). On the other hand, under complete semantics, the query $\exists x. \text{author}(x, \text{john})$, asking if *john* is author of some argument, is true under credulous reasoning and false under skeptical reasoning (as *john*'s argument a belongs to E_2 only).

Consider now the case where there are two people from France and UK and they are interested in opinions expressed by people resident in their respective countries. That is, the subframeworks of interest are those where only opinions of French and UK people, respectively, are considered. This can be achieved by extending KAF with constraints expressed by means of restricted relational calculus formulae (Abiteboul, Hull, and Vianu 1995). The safe, range restricted relational calculus provides adequate power for many applications and at the same time can be implemented efficiently.

Example 4. The KAF subframework containing only opinions from UK people can be selected by using the following constraint: $\forall x. \text{arg}(x) \Rightarrow \exists y, z. \text{author}(x, y) \wedge \text{resident}(y, z) \wedge \text{nation}(z, \text{uk})$. $\square$

The *Constrained* KAF (CKAF) is a KAF where a set of restricted relational calculus formulae is used for reasoning over the subframeworks that satisfy the formulae and *minimally* differ from the original framework. Two minimality criteria are considered, *set* and *card* minimality. Hence, a subframework is obtained by deleting a minimal set or a minimal number of arguments and attacks, as well as the related information in the KB, so that the set of constraints is satisfied w.r.t. the resulting KAF.

Intuitively, a CKAF concisely represents a set of KAFs that could be obtained by extracting ‘feasible’ subframeworks. Then, its semantics can be defined in terms of the semantics of its subframeworks. This means that, similarly to other frameworks previously introduced in the literature (e.g., iAF, PrAF), the subframeworks can be seen as possible worlds. Thus, classical verification and acceptance problems (i.e., credulous and skeptical acceptance) need to be reconsidered in light of the presence of possible worlds, leading to the concepts of *possible* and *necessary verification*, as well as *possible* and *necessary* credulous/skeptical acceptance (cf. Definitions 11 and 12).

Contributions. Our main contributions are as follows.¹

- We introduce KAF, a framework where AF is complemented with a *DL-Lite* KB containing additional information related to arguments. KAF extensions are defined in terms of argumentation semantics and consists of *DL-Lite* KBs, over which union of conjunctive queries are posed.
- We show that the complexity of acceptance problems in KAF does not increase with respect to that of AF.
- We introduce CKAF, which extends KAF with constraints, expressed by means of restricted relational calculus formulae, allowing reasoning over selected subframeworks (i.e., KAFs) satisfying the formula and differing from the original framework by a minimal set (or a minimal number) of elements, in terms of arguments and attacks.
- We introduce the verification and acceptance problems for CKAF and thoroughly investigate the complexity of these problems. A summary of our results for the verification and possible/necessary credulous/skeptical acceptance under semantics $\sigma \in \{\text{gr}, \text{co}, \text{pr}, \text{st}\}$ and under set and cardinality subframework semantics are reported in Table 1 and Table 2, respectively. It turns out that those problems in CKAF are harder than corresponding ones in iAF (Baumeister et al. 2021), correlated iAF (Fazzinga, Flesca, and Furfaro 2021) and constrained iAF (Mailly 2024).
- Consistently with our complexity results, we show that iAF, correlated and constrained iAF can be viewed as special cases of CKAF, in the sense that each of them can be rewritten into an extensions-equivalent CKAF, under both set and cardinality subframework semantics.

Integrating AF with DLs addresses a crucial need in knowledge representation: managing structured reasoning under uncertainty and conflicts in expressive, ontology-based systems. While DLs provide a formal foundation for representing and reasoning about conceptual knowledge

in a decidable and often tractable way, they are inherently monotonic, meaning they cannot naturally handle conflicting or defeasible information. In contrast, AF offers a non-monotonic reasoning mechanism that evaluates conflicting arguments and supports decision-making based on acceptability semantics. Combining these paradigms leverages the strengths of both: DLs handle structured domain knowledge, such as medical ontologies, legal codes, or taxonomies, while AFs manage conflicts, priorities, and justifications among inferences derived from this knowledge. For example, in medical diagnosis, DLs can encode clinical guidelines and taxonomies of symptoms and conditions, whereas argumentation can evaluate competing diagnostic hypotheses, especially when guidelines or observations conflict.

2 Preliminaries

In this section, after briefly recalling the Dung’s framework, we review basic concepts underlying Description Logics.

2.1 Argumentation Framework

An abstract *Argumentation Framework* (AF) (Dung 1995) is a pair $\langle A, R \rangle$, where A is a set of *arguments* and $R \subseteq A \times A$ is a set of *attacks*.

Given an AF $\Lambda = \langle A, R \rangle$ and a set $S \subseteq A$ of arguments, an argument $a \in A$ is said to be *i) defeated* w.r.t. S iff there exists $b \in S$ such that $(b, a) \in R$, and *ii) acceptable* w.r.t. S iff for every argument $b \in A$ with $(b, a) \in R$, there is $c \in S$ such that $(c, b) \in R$. The sets of arguments defeated and acceptable w.r.t. S are as follows (where Λ is fixed):

- $\text{Def}(S) = \{a \in A \mid \exists b \in S. (b, a) \in R\}$;
- $\text{Acc}(S) = \{a \in A \mid \forall b \in A. (b, a) \in R \Rightarrow b \in \text{Def}(S)\}$.

Given an AF $\langle A, R \rangle$, a set $S \subseteq A$ of arguments is said to be *conflict-free* iff $S \cap \text{Def}(S) = \emptyset$. Moreover, $S \subseteq A$ is said to be a *complete* (co) extension iff it is conflict-free and $S = \text{Acc}(S)$. A complete extension $S \subseteq A$ is said to be:

- *preferred* (pr) iff it is $\subseteq$ -maximal;
- *stable* (st) iff it is a total preferred extension, that is, a preferred extension such that $S \cup \text{Def}(S) = A$;
- *grounded* (gr) iff it is $\subseteq$ -minimal.

In the following, if not specified otherwise, σ denotes any semantics in $\{\text{gr}, \text{co}, \text{st}, \text{pr}\}$. For any AF Λ and semantics σ , $\sigma(\Lambda)$ denotes the set of σ -extensions of Λ . All the above-mentioned semantics except the stable admit at least one extension (i.e., $\sigma(\Lambda) \neq \emptyset$), and the grounded admits exactly one extension (i.e., $|\text{gr}(\Lambda)| = 1$) (Dung 1995; Caminada 2006). For any AF Λ , it holds that $\text{st}(\Lambda) \subseteq \text{pr}(\Lambda) \subseteq \text{co}(\Lambda)$, and $\text{gr}(\Lambda) \subseteq \text{co}(\Lambda)$.

For any AF $\Lambda = \langle A, R \rangle$, semantics σ , and argument $a \in A$, we say that a is *credulously* (resp. *skeptically*) accepted (under semantics σ), denoted as $\text{CA}_\sigma(a, \Lambda)$ (resp. $\text{SA}_\sigma(a, \Lambda)$) if a belongs to at least one (resp. every) σ -extension of Λ . Moreover, we use CA_σ (resp. SA_σ), or CA (resp. SA) whenever σ is understood, to denote the credulous (resp. skeptical) acceptance problem, that is, the problem of deciding whether an argument is credulously (resp. skeptically) accepted. Clearly, for the grounded semantics,

¹We assume the reader is familiar with basic concepts of computational complexity.

which has exactly one extension, these problems are identical (i.e., $CA_{gr} \equiv SA_{gr}$).

Example 5. Consider the AF $\Lambda = \langle \{a, b, c, d\}, \{(a, b), (b, a), (a, c), (b, c), (c, d), (d, c)\} \rangle$. Λ has 4 complete extensions: $E_0 = \emptyset$, $E_1 = \{a, d\}$, $E_2 = \{b, d\}$ and $E_3 = \{d\}$. That is, $co(\Lambda) = \{E_0, E_1, E_2, E_3\}$. E_0 is the grounded extension, E_1 and E_2 are stable (and preferred) extensions. Thus, a, b, d are credulously accepted for all semantics except for the grounded; only d is skeptically accepted under stable and preferred semantics. $\square$

2.2 Description Logics

Syntax. A DL *knowledge base* (KB), hereafter simply called a knowledge base, consists of an *Assertional Box* ABox (database) and a *Terminological Box* TBox (ontology), which are constructed from a set N_C of *concept names* (unary predicates), a set of N_R of *role names* (binary predicates), and a set N_I of *individuals* (constants). Specifically, an ABox is a finite set of *concept assertions* of the form $A(a)$, where $A \in N_C$ and $a \in N_I$, and *role assertions* of the form $R(a, b)$, where $R \in N_R$ and $a, b \in N_I$. A TBox is a finite set of axioms whose form depends on the considered DL. In *DL-Lite_{core}*, TBox axioms are only *concept inclusions* of the form $B \sqsubseteq C$ formed using the following syntax (where $A \in N_C$ and $R \in N_R$):

$$\begin{aligned} B &:= A \mid \exists Q; & Q &:= R \mid R^-; \\ C &:= B \mid \neg B; & S &:= Q \mid \neg Q. \end{aligned}$$

Several extensions of *DL-Lite_{core}* have been proposed in the literature, including *DL-Lite_R*, which additionally allows *role inclusions* of the form $Q \sqsubseteq S$, and *DL-Lite_F*, which allows expressing functional dependencies on roles. In this paper, we consider *DL-Lite_R*, which provides the logical underpinnings of the OWL 2 QL profile (W3C 2012) and whenever we use the term *DL-Lite* we mean *DL-Lite_R*, although our results hold for all (query rewritable) *DL-Lite* fragments (Calvanese et al. 2007; Cima, Lenzerini, and Poggi 2020; Bienvenu 2016; Artale et al. 2009; Bienvenu and Ortiz 2015). For a given TBox, we partition its (concept or role) names into *derived* ones, which appear on the right-hand side of (concept or role) inclusions in the TBox, and *base* ones, which are the remaining names. Such terminology is borrowed from the deductive databases field, where predicates are partitioned into base and derived.

Semantics. An *interpretation* is of the form $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, where $\Delta^{\mathcal{I}}$ is a non-empty set called *interpretation domain* and $\cdot^{\mathcal{I}}$ is an *interpretation function* mapping each $A \in N_C$ to $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$, each $R \in N_R$ to $R^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$, and each $a \in N_I$ to $a^{\mathcal{I}} \in \Delta^{\mathcal{I}}$, with $a \neq b$ implying $a^{\mathcal{I}} \neq b^{\mathcal{I}}$ (the last condition is the well-known Unique Names Assumption). The function $\cdot^{\mathcal{I}}$ is extended to general concepts and roles as follows: $(\exists Q)^{\mathcal{I}} = \{c \mid \exists d. (c, d) \in Q^{\mathcal{I}}\}$, $(\neg B)^{\mathcal{I}} = \Delta^{\mathcal{I}} \setminus B^{\mathcal{I}}$, $(R^-)^{\mathcal{I}} = \{(c, d) \mid (d, c) \in R^{\mathcal{I}}\}$, $(\neg Q)^{\mathcal{I}} = (\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}) \setminus Q^{\mathcal{I}}$. An interpretation $\mathcal{I}$ satisfies an inclusion $G \sqsubseteq H$ if $G^{\mathcal{I}} \subseteq H^{\mathcal{I}}$; it satisfies $A(a)$ (resp. $R(a, b)$) if $a^{\mathcal{I}} \in A^{\mathcal{I}}$ (resp. $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in R^{\mathcal{I}}$). An interpretation $\mathcal{I}$ is a *model* of a KB $\langle \mathcal{T}, \mathcal{A} \rangle$ iff $\mathcal{I}$ satisfies all inclusions in $\mathcal{T}$ and all assertions in $\mathcal{A}$. A KB is consistent if it has a model, otherwise it is inconsistent.

Querying. A *conjunctive query* (CQ) is an expression of the form: $q(\vec{x}) := \exists \vec{y}. \varphi(\vec{x}, \vec{y})$, where $\varphi(\vec{x}, \vec{y})$ is a conjunction of atoms of the forms $A(t)$ or $R(t, t')$, where $A \in N_C$, $R \in N_R$, t, t' are variables or individuals from N_I , $\vec{x}$ denotes the set of free variables, $\vec{y}$ denotes the set of existentially quantified variables. A CQ is Boolean if all of its variables are existentially quantified (and thus $\vec{x}$ is empty); in such a case the query is denoted as $q()$ or simply q . A Boolean CQ q is entailed by a knowledge base $\mathcal{K}$, denoted $\mathcal{K} \models q$, iff q holds in all models of $\mathcal{K}$. A (*certain*) *answer* to a non-Boolean CQ $q(\vec{x})$, where $\vec{x} = (x_1, \dots, x_k)$, w.r.t. a knowledge base $\mathcal{K}$ is a tuple of individuals $\vec{c} = (c_1, \dots, c_k)$ such that $\mathcal{K} \models q(\vec{c})$, where $q(\vec{c})$ is the Boolean CQ query obtained from $q(\vec{x})$ by replacing each x_i by c_i . We use $Ans(q(\vec{x}), \mathcal{K})$ to denote the set of all answers to $q(\vec{x})$ w.r.t. $\mathcal{K}$.

A union of conjunctive queries (UCQ) allows the disjunctions of multiple CQs and is an expression of the form:

$$q(\vec{x}) := q_1(\vec{x}) \vee \dots \vee q_n(\vec{x}),$$

where each $q_i(\vec{x})$ is a CQ. A UCQ q is Boolean if all of its variables are existentially quantified, and it is entailed by a knowledge base $\mathcal{K}$, denoted $\mathcal{K} \models q$, iff q holds in all models of $\mathcal{K}$. The set of all answers to a non-Boolean UCQ $q(\vec{x})$ w.r.t. a knowledge base $\mathcal{K}$ is $Ans(q(\vec{x}), \mathcal{K}) = \bigcup_{i=1}^n Ans(q_i(\vec{x}), \mathcal{K})$.

To express constraints (introduced in Section 4) we will use the safe, range restricted relational calculus; we refer to (Abiteboul, Hull, and Vianu 1995) for details on it.

3 Knowledge-based AF

We now introduce an extension of Dung's AF that consists in combining an AF with a knowledge base (defined in *DL-Lite*). The derived framework is called *Knowledge-based Argumentation Framework* (KAF).

Syntax. From now on, we will assume that N_I contains a set N_A of distinguished individuals used to refer to arguments. We will also assume the existence of a built-in concept name *arg*, taking values from N_A , and a built-in role name called *att*, taking values from $N_A \times N_A$. We also impose that for any TBox, *arg* and *att* are base names (cf. Section 2.2). For an ABox $\mathcal{A}$, we define $Arg(\mathcal{A}) = \{a \mid arg(a) \in \mathcal{A}\}$ and $Att(\mathcal{A}) = \{(a, b) \mid att(a, b) \in \mathcal{A}\}$.

Definition 1 (KAF). A Knowledge-based AF (KAF) is a quadruple $\Omega = \langle A, R, \mathcal{T}, \mathcal{A} \rangle$ where $\langle A, R \rangle$ is an AF and $\langle \mathcal{T}, \mathcal{A} \rangle$ is a consistent KB s.t. $Arg(\mathcal{A}) = A$ and $Att(\mathcal{A}) = R$.

Thus, arguments and attacks are included by default in the ABox $\mathcal{A}$, whereas the limitation to the form of assertions states that arguments and attacks (referring to the AF component) cannot be defined by means of inclusions in $\mathcal{T}$.

Example 6. The scenario of Example 1 can be encoded in the KAF $\Omega = \langle A, R, \mathcal{T}, \mathcal{A} \rangle$, where:

$$\begin{aligned} A &= \{a, b, c, d, e\}; \\ R &= \{(a, b), (b, a), (c, b), (e, c), (c, d), (d, c)\}; \\ \mathcal{T} &= \{(arg \sqsubseteq \exists author), (\exists author \sqsubseteq \exists resident)\}; \end{aligned}$$

$$\mathcal{A} = \{arg(x) \mid x \in A\} \cup \{att(x, y) \mid (x, y) \in R\} \cup$$

$$\left\{ \begin{array}{ll} \text{author}(\text{a}, \text{john}), & \text{resident}(\text{john}, \text{paris}), \\ \text{author}(\text{b}, \text{mary}), & \text{resident}(\text{mary}, \text{paris}), \\ \text{author}(\text{c}, \text{frank}), & \text{resident}(\text{frank}, \text{bath}), \\ \text{author}(\text{d}, \text{anne}), & \text{resident}(\text{anne}, \text{bath}), \\ \text{nation}(\text{paris}, \text{france}), & \text{nation}(\text{bath}, \text{uk}) \end{array} \right\}$$

The assertions of KAF Ω are illustrated in Figure 1. $\square$

Given a KAF $\Omega = \langle A, R, \mathcal{T}, \mathcal{A} \rangle$ and a set of arguments $S \subseteq A$, we denote by $R_{\downarrow S} = R \cap (S \times S)$ and with $\mathcal{A}_{\downarrow S} = \mathcal{A} \setminus (\{p(a) \in \mathcal{A} \mid a \in A \setminus S\} \cup \{p(a, b) \in \mathcal{A} \mid \{a, b\} \cap (A \setminus S) \neq \emptyset\})$ the projections over S of R and $\mathcal{A}$, respectively. Moreover, we denote with $\Omega_{\downarrow S} = \langle S, R_{\downarrow S}, \mathcal{T}, \mathcal{A}_{\downarrow S} \rangle$ the projection of Ω over S . Finally, we define $af(A, R) = \{arg(a) \mid a \in A\} \cup \{att(a, b) \mid (a, b) \in R\}$ as the ABox consisting of the assertions that describe the AF $\langle A, R \rangle$. We next define the concept of subframework, that is the projection of a KAF over a set of arguments, that will be used in Section 4 for defining CKAF semantics.

Definition 2 (Subframework). Given a KAF $\Omega = \langle A, R, \mathcal{T}, \mathcal{A} \rangle$ and a set of arguments $S \subseteq A$, we say that $\Omega_{\downarrow S} = \langle S, R_{\downarrow S}, \mathcal{T}, \mathcal{A}_{\downarrow S} \rangle$ is a *subframework* of Ω .

Example 7. Consider the KAF $\Omega = \langle A, R, \mathcal{T}, \mathcal{A} \rangle$ of Example 6, and the set $S = \{a, d, e\}$. We have that $\Omega_{\downarrow S} = \langle \{a, d, e\}, \emptyset, \mathcal{T}, \mathcal{A} \setminus (\{arg(b), arg(c)\} \cup \{att(x, y) \mid (x, y) \in R\} \cup \{author(b, mary), author(c, frank)\}) \rangle$. $\square$

Semantics. We first extend the concept of extension previously defined for AF, and then consider query acceptance.

Definition 3 (KAF Semantics). Given a KAF $\Omega = \langle A, R, \mathcal{T}, \mathcal{A} \rangle$ and an argumentation semantics $\sigma \in \{\text{gr}, \text{co}, \text{st}, \text{pr}\}$, a σ -extension for Ω is a KB $\langle \mathcal{T}, \mathcal{A}_{\downarrow E} \rangle$ such that $E \in \sigma(\langle A, R \rangle)$.

The set of σ -extensions for a KAF Ω is denoted by $\sigma(\Omega)$.

As an example, for the KAF Ω of Example 6 and $S = \{a, d, e\}$ of Example 7, $\mathcal{K} = \langle \mathcal{T}, \mathcal{A}_{\downarrow S} \rangle$ is a stable (and also preferred) extension, that is $\mathcal{K} \in \text{st}(\Omega)$.

Querying. In the following, we use the letter α to denote the type of acceptance, that is $\alpha = cr$ means that we refer to credulous acceptance, whereas $\alpha = sk$ means that we are referring to the skeptical acceptance.

Definition 4 (Query answer). Let $\Omega = \langle A, R, \mathcal{T}, \mathcal{A} \rangle$ be a KAF, $\sigma \in \{\text{gr}, \text{co}, \text{st}, \text{pr}\}$ an argumentation semantics, $q(\vec{x})$ a UCQ, and $\alpha \in \{cr, sk\}$ a type of acceptance. Then, the set of answers to the query $q(\vec{x})$ over Ω under semantics σ and acceptance α , denoted by $Ans_{\sigma}^{\alpha}(q(\vec{x}), \Omega)$, is:

- $Ans_{\sigma}^{cr}(q(\vec{x}), \Omega) = \{\vec{a} \mid \exists \mathcal{K} \in \sigma(\Omega). \vec{a} \in Ans(q(\vec{x}), \mathcal{K})\};$
- $Ans_{\sigma}^{sk}(q(\vec{x}), \Omega) = \{\vec{a} \mid \forall \mathcal{K} \in \sigma(\Omega). \vec{a} \in Ans(q(\vec{x}), \mathcal{K})\}.$

Definition 5 (Acceptance). For any KAF $\Omega = \langle A, R, \mathcal{T}, \mathcal{A} \rangle$, argumentation semantics $\sigma \in \{\text{gr}, \text{co}, \text{st}, \text{pr}\}$, tuple of individuals $\vec{a}$, and UCQ $q(\vec{x})$, we say that $\vec{a}$ is:

- *credulously accepted* if $\vec{a} \in Ans_{\sigma}^{cr}(q(\vec{x}), \Omega)$;
- *skeptically accepted* if $\vec{a} \in Ans_{\sigma}^{sk}(q(\vec{x}), \Omega)$.

We use CA_{σ} (resp. SA_{σ}), to denote the credulous (resp. skeptical) acceptance problem in KAF under semantics σ , that is, the problem of deciding whether a tuple of individuals is credulously (resp. skeptically) accepted.

Example 8. Consider the KAF $\Omega = \langle A, R, \mathcal{T}, \mathcal{A} \rangle$ of Example 6, and the query $q(x) = \text{author}(x, \text{john})$, asking for the arguments authored by john. Under stable semantics, we have that $Ans_{\text{st}}^{cr}(q(\vec{x}), \Omega) = \{a\}$ and $Ans_{\text{st}}^{sk}(q(\vec{x}), \Omega) = \emptyset$. Thus, under stable semantics, a is credulously accepted, but not skeptically accepted. $\square$

Proposition 1. For KAF, it holds that:

- CA_{σ} is *i*) in *PTIME* for $\sigma = \text{gr}$, and *ii*) *NP-complete* for $\sigma \in \{\text{co}, \text{st}, \text{pr}\}$;
- SA_{σ} is *i*) in *PTIME* for $\sigma \in \{\text{gr}, \text{co}\}$, *ii*) *coNP-complete* for $\sigma = \text{st}$, and *iii*) Π_2^p -complete for $\sigma = \text{pr}$.

Therefore, the complexity of acceptance problems in KAF and AF coincides. That is, from the computational standpoint, adding the auxiliary information in the DL KB comes at no additional cost. We will show that this is not the case of CKAF, formally introduced in the next section.

4 Constrained Knowledge-based AF

We now introduce the *Constrained Knowledge-based AF* (CKAF), that extends KAF with constraints expressed by means of restricted relational calculus formulae.

Syntax. Constraints are introduced to pose restrictions over a CKAF framework and are used to define the subframeworks the user is interested in.

Definition 6 (CKAF). A Constrained Knowledge-based AF (CKAF) is a quintuple $\Delta = \langle A, R, \mathcal{T}, \mathcal{A}, \Gamma \rangle$ where $\Omega_{\Delta} = \langle A, R, \mathcal{T}, \mathcal{A} \rangle$ is a KAF and Γ is a set of logic formulae (denoting constraints) of the form $\forall \vec{x}. \psi(\vec{x}) \Rightarrow \varphi(\vec{x})$, where $\psi(\vec{x})$ and $\varphi(\vec{x})$ are either *i*) a safe, range restricted relational query using base (concept or role) names, or *ii*) a union of conjunctive queries using arbitrary (concept or role) names.

Clearly, variables in $\vec{x}$ occur as free variables in $\psi(\vec{x})$ and $\varphi(\vec{x})$. Moreover, when $\vec{x}$ is empty the two queries are Boolean. To make more explicit our constraints, we sometimes write constraints of the form $\text{true} \Rightarrow \varphi()$ (resp. $\psi() \Rightarrow \text{false}$), where *true* (resp. *false*) stands for a Boolean query always returning the truth value *true* (resp. *false*),² which intuitively expresses that $\varphi()$ must be true (resp. $\psi()$ must be false). In the following, whenever Γ consists of a single formula f , we will write $\Gamma = f$ instead of $\Gamma = \{f\}$.

Notice that the constraints in Γ are trivially satisfied whenever the sets of base concepts and roles are empty.

Semantics. As said above, Γ contains constraints posing restrictions on the form of Ω_{Δ} . We start by defining the satisfaction of the formulae in Γ with respect to the underlying KAF framework. A constraint $\forall \vec{x}. \psi(\vec{x}) \Rightarrow \varphi(\vec{x})$ is satisfied by a KB $\mathcal{K}$ iff $Ans(\psi(\vec{x}), \mathcal{K}) \subseteq Ans(\varphi(\vec{x}), \mathcal{K})$. Γ is satisfied by $\mathcal{K}$ iff every constraint in Γ is satisfied by $\mathcal{K}$. A KAF $\Omega = \langle A, R, \mathcal{T}, \mathcal{A} \rangle$ satisfies a set of constraints Γ , denoted as $\Omega \models \Gamma$, iff $\langle \mathcal{T}, \mathcal{A} \rangle \models \Gamma$.

As we are considering UCQs and *DL-Lite* knowledge bases, as well as safe, range restricted relational calculus

²Observe that it does not make sense to write constraints with the truth value *false* in the body of implication and with the truth value *true* in the head.

queries over the ABox, query answers are computable in AC_0 and are finite (Calvanese et al. 2007; Cima, Lenzerini, and Poggi 2020; Bienvenu 2016). The role of the logical formulae Γ in Δ is to express constraints to be satisfied by the underlying Ω_Δ . Intuitively, if some of the logical formulae in Γ is not satisfied, then the underlying KAF should be revised by computing subframeworks, that is, by (minimally) modifying the topology of the underlying AF through the deletion of arguments, so that the formula is satisfied. Clearly, when deleting arguments, incident roles (and, thus, also attacks) are deleted as well.

Given a preorder $\preceq$ over all subframeworks of a KAF Ω , a subframework $\Omega' = \langle A', R', \mathcal{T}, \mathcal{A}' \rangle$ of Ω is $\preceq$ -maximal iff there is no subframework $\Omega'' = \langle A'', R'', \mathcal{T}, \mathcal{A}'' \rangle$ of Ω such that $\Omega' \prec \Omega''$. Hereafter, we consider two classical preorders (over the subframeworks of the same KAF):

- $\Omega' \preceq^c \Omega''$ iff $|A'| + |R'| \leq |A''| + |R''|$;
- $\Omega' \preceq^s \Omega''$ iff $A' \cup R' \subseteq A'' \cup R''$.

Thus, $\preceq^s$ compares subframeworks on the basis of the set of arguments and attacks, whereas $\preceq^c$ compares them by focusing on the number of arguments and attacks.

Definition 7 ($\preceq$ -Subframework). *A $\preceq$ -subframework for a CKAF $\Delta = \langle A, R, \mathcal{T}, \mathcal{A}, \Gamma \rangle$ is a subframework Ω' of Ω_Δ such that i) $\Omega' \models \Gamma$ and ii) there is no subframework Ω'' of Ω_Δ such that $\Omega'' \models \Gamma$ and $\Omega' \prec \Omega''$.*

Thus, a $\preceq^s$ -subframework (*s-subframework* for short) is obtained by deleting a minimal (w.r.t. set-inclusion) set of arguments and attacks, whereas a $\preceq^c$ -subframework (*c-subframework*) is obtained by deleting a minimum number of arguments and attacks.

Given a preorder $\preceq^\delta$ with $\delta \in \{s, c\}$, the set of δ -subframeworks for a CKAF Δ is denoted by $\mathcal{S}^\delta(\Delta)$. The semantics of Δ is given by the set of δ -subframeworks of Δ .

Example 9. Consider the CKAF $\Delta = \langle A, R, \mathcal{T}, \mathcal{A}, \Gamma \rangle$ where $\Omega = \langle A, R, \mathcal{T}, \mathcal{A} \rangle$ is the KAF of Example 6, and $\Gamma = \forall x. \arg(x) \Rightarrow \exists y, z. \text{author}(x, y) \wedge \text{resident}(y, z) \wedge \text{nation}(z, \text{uk})$ is the formula of Example 4 whose aim is to restrict to opinions expressed by people resident in the UK only. The only δ -subframework for Δ , with $\delta \in \{s, c\}$, is $\mathcal{S}^\delta(\Delta) = \{\Omega_{\downarrow\{c, d\}}\}$. $\square$

Querying. To answer queries we must consider the subframeworks of interest and, therefore, the type of preorder.

Definition 8 (Query answer). *Let $\Delta = \langle A, R, \mathcal{T}, \mathcal{A}, \Gamma \rangle$ be a CKAF, $\sigma \in \{\text{gr}, \text{co}, \text{st}, \text{pr}\}$ an argumentation semantics, $q(\vec{x})$ a UCQ, $\alpha \in \{cr, sk\}$ the type of acceptance, and $\delta \in \{s, c\}$ a subframework semantics. Then, the set of possible (resp. necessary) δ -subframework answers to the query $q(\vec{x})$ over Δ under semantics σ and acceptance type α is:*

- (possible-subframework) $\delta\text{-Ans}_\sigma^{\text{p}, \alpha}(q(\vec{x}), \Delta) = \{\vec{a} \mid \exists \Omega \in \mathcal{S}^\delta(\Delta). \vec{a} \in \text{Ans}_\sigma^\alpha(q(\vec{x}), \Omega)\};$
- (necessary-subframework) $\delta\text{-Ans}_\sigma^{\text{n}, \alpha}(q(\vec{x}), \Delta) = \{\vec{a} \mid \forall \Omega \in \mathcal{S}^\delta(\Delta). \vec{a} \in \text{Ans}_\sigma^\alpha(q(\vec{x}), \Omega)\}.$

Definition 9 (Acceptance). *Let $\Delta = \langle A, R, \mathcal{T}, \mathcal{A}, \Gamma \rangle$ be a CKAF, $\sigma \in \{\text{gr}, \text{co}, \text{st}, \text{pr}\}$ an argumentation semantics, $q(\vec{x})$ a UCQ, $\alpha \in \{cr, sk\}$ the type of acceptance, $\delta \in \{s, c\}$ a subframework semantics, and $\vec{a}$ tuple of individuals. We say that $\vec{a}$ is:*

- *possibly credulously accepted* if $\vec{a} \in \delta\text{-Ans}_\sigma^{\text{p}, cr}(q(\vec{x}), \Delta)$;
- *possibly skeptically accepted* if $\vec{a} \in \delta\text{-Ans}_\sigma^{\text{p}, sk}(q(\vec{x}), \Delta)$;
- *necessary credulously accepted* if $\vec{a} \in \delta\text{-Ans}_\sigma^{\text{n}, cr}(q(\vec{x}), \Delta)$;
- *necessary skeptically accepted* if $\vec{a} \in \delta\text{-Ans}_\sigma^{\text{n}, sk}(q(\vec{x}), \Delta)$.

Example 10. Consider the CKAF $\Delta' = \langle A, R, \emptyset, \mathcal{A}, \Gamma' \rangle$ derived from the CKAF Δ of Example 9 by replacing i) $\mathcal{T}$ with an empty set of axioms (so that author and resident become base roles), and ii) Γ with $\Gamma' = \{\forall x_1, x_2. \exists y_1, z_1, y_2, z_2. \text{author}(x_1, y_1) \wedge \text{resident}(y_1, z_1) \wedge \text{author}(x_2, y_2) \wedge \text{resident}(y_2, z_2) \Rightarrow z_1 = z_2\}$.³ Let $\Omega = \langle A, R, \emptyset, \mathcal{A} \rangle$, for Δ' there are two s-sub-frameworks: $\Omega_1 = \Omega_{\downarrow\{a, b, e\}}$ (containing the arguments whose authors reside in paris or not authored by anyone), and $\Omega_2 = \Omega_{\downarrow\{c, d, e\}}$ (containing only the arguments whose authors reside in bath or not authored by anyone). Only Ω_2 is also a c-subframework, as it is obtained by deleting two arguments and three attacks, whereas to obtain Ω_1 two arguments and four attacks must be deleted.

Take now the query $q_1(x, y) = \text{author}(x, y)$ over Δ' , under $\preceq^c$ -preorder. As Ω_2 is the only c-subframework for Δ' , whose underlying AF has exactly one complete extension $\{e, d\}$, under all semantics $\sigma \in \{\text{gr}, \text{co}, \text{st}, \text{pr}\}$ and acceptance type $\alpha \in \{cr, sk\}$ we have that:

$$c\text{-Ans}_\sigma^{\text{p}, \alpha}(q_1(x, y), \Delta') = c\text{-Ans}_\sigma^{\text{n}, \alpha}(q_1(x, y), \Delta') = \{\langle d, \text{anne} \rangle\}.$$

Taking now the $\preceq^s$ preorder, both Ω_1 and Ω_2 are s-subframeworks of Δ' . Under stable semantics, Ω_1 has two extensions, whereas Ω_2 has only one extension. Therefore:

- $s\text{-Ans}_{\text{st}}^{\text{p}, cr}(q_1(x, y), \Delta') = \{\langle a, \text{john} \rangle, \langle b, \text{mary} \rangle, \langle d, \text{anne} \rangle\};$
- $s\text{-Ans}_{\text{st}}^{\text{p}, sk}(q_1(x, y), \Delta') = \{\langle d, \text{anne} \rangle\};$
- $s\text{-Ans}_{\text{st}}^{\text{n}, cr}(q_1(x, y), \Delta') = s\text{-Ans}_{\text{st}}^{\text{n}, sk}(q_1(x, y), \Delta') = \{\}. \quad \square$

Note that a CKAF without base predicates is similar to a correlated/constrained incomplete AF (Fazzinga, Flesca, and Furfaro 2020; Mailly 2020), though the formulae defined in those works are propositional and allow only built-in predicates. Moreover, even the aim is different, as we are interested in dealing with subframeworks (which minimally differ from the original framework), whereas in the above-mentioned works all (valid) *completions* are first-class citizens, that is, they do not rely on minimal changes.

5 Computational Complexity

We investigate the complexity of fundamental reasoning problems for CKAF. In particular, we study the verification, existence, and credulous/skeptical acceptance problems, that are usually considered for analyzing the complexity of argumentation frameworks. Our complexity results refer to data complexity, where the TBox component, the constraints and the query are assumed to be fixed.

We start with a lemma stating that the satisfaction of the logical formulae in Γ can be checked in PTIME in the size of the active domain consisting of constant values.

³Although Γ' is not formulated according to the syntax of Definition 6, it can be rewritten in a standard Boolean relational calculus query by moving the head of the implication into the body and rewriting the universal quantifier by negating the formula twice.

	AF		iAF		ciAF		CKAF			
σ	V_σ	PV_σ	NV_σ	PV_σ	NV_σ		$c\text{-}PV_\sigma$	$s\text{-}PV_\sigma$	$c\text{-}NV_\sigma$	$s\text{-}NV_\sigma$
gr	P	P	P	NP-c	coNP-c		$\Theta_2^p\text{-c}$	$\Sigma_2^p\text{-c}$	$\Theta_2^p\text{-c}$	$\Pi_2^p\text{-c}$
co	P	P	P	NP-c	coNP-c		$\Theta_2^p\text{-c}$	$\Sigma_2^p\text{-c}$	$\Theta_2^p\text{-c}$	$\Pi_2^p\text{-c}$
st	P	P	P	NP-c	coNP-c		$\Theta_2^p\text{-c}$	$\Sigma_2^p\text{-c}$	$\Theta_2^p\text{-c}$	$\Pi_2^p\text{-c}$
pr	coNP-c	$\Sigma_2^p\text{-c}$	coNP-c	$\Sigma_2^p\text{-c}$	coNP-c		$\Theta_2^p\text{-h}, \Sigma_2^p$	$\Sigma_2^p\text{-c}$	$\Theta_2^p\text{-c}$	$\Pi_2^p\text{-c}$

Table 1: Complexity of the verification problems for AF, iAF, ciAF, and CKAF under semantics $\sigma \in \{\text{gr}, \text{co}, \text{pr}, \text{st}\}$. For any complexity class C , $C\text{-c}$ (resp. $C\text{-h}$) means C -complete (resp. C -hard). An interval $C\text{-h}$, C' means C -hard and in C' .

Lemma 1. Given a CKAF $\Delta = \langle A, R, \mathcal{T}, \mathcal{A}, \Gamma \rangle$ and a subframework $\Omega' = \langle A', R', \mathcal{T}, \mathcal{A}' \rangle$ of $\Omega = \langle A, R, \mathcal{T}, \mathcal{A} \rangle$, checking whether $\Omega' \models \Gamma$ can be done in PTIME.

Observe that a given CKAF may admit zero, one, or multiple KAF subframeworks. Subframework existence, verification, and acceptance problems can be defined analogously to those defined for iAF (Baumeister et al. 2018; Baumeister et al. 2021). In fact, roughly speaking, subframeworks correspond to iAF completions which, however, must satisfy a minimality criteria (either set or cardinality subframework semantics).

5.1 Existence Problems

We start by introducing the *existence problems* under set and cardinality subframework semantics, i.e., $\delta \in \{c, s\}$.

Definition 10. Let $\delta \in \{c, s\}$ be a subframework semantics, $\delta\text{-EX}$ is the problem of checking whether there exists a δ -subframework for a given CKAF.

In our running example, both problems are true. However, deciding subframework existence is hard.

Theorem 1. $s\text{-EX}$ and $c\text{-EX}$ are NP-complete.

5.2 Verification Problems

We now characterize the complexity of the *verification problems* for CKAF, which are formally defined in what follows.

Definition 11. Let $\Delta = \langle A, R, \mathcal{T}, \mathcal{A}, \Gamma \rangle$ be a CKAF, $\mathcal{A}^* \subseteq \mathcal{A}$, $\sigma \in \{\text{gr}, \text{co}, \text{pr}, \text{st}\}$ an argumentation semantics, and $\delta \in \{c, s\}$ a subframework semantics. Then:

- the δ -possible verification problem, denoted as $\delta\text{-PV}_\sigma$, is the problem of checking whether $\langle \mathcal{T}, \mathcal{A}^* \rangle$ is a σ -extension in any δ -subframework for Δ ;
- the δ -necessary verification problem, denoted as $\delta\text{-NV}_\sigma$, is the problem of checking whether $\langle \mathcal{T}, \mathcal{A}^* \rangle$ is a σ -extension in all δ -subframeworks for Δ .

Observe that for a CKAF of the form $\langle A, R, \emptyset, af(A, R), \emptyset \rangle$, that semantically coincides with the AF $\langle A, R \rangle$, $s\text{-PV}_\sigma = c\text{-PV}_\sigma = s\text{-NV}_\sigma = c\text{-NV}_\sigma$, which in turns coincide with the standard verification problem for AF, that is checking whether a given set $S \subseteq A$ of arguments is a σ -extension of AF $\langle A, R \rangle$.

The following theorem states that the complexity of the verification problem for CKAF is generally harder than that

for AF, whose complexity is summarized in the second column of Table 1 (Dvorak and Dunne 2017). Moreover, the results for the cardinality-based semantics are generally lower than those for the set-based semantics.

Theorem 2.

- $c\text{-PV}_\sigma$ is i) Θ_2^p -complete for $\sigma \in \{\text{gr}, \text{co}, \text{st}\}$, and ii) Θ_2^p -hard and in Σ_2^p for $\sigma = \text{pr}$;
- $s\text{-PV}_\sigma$ is Σ_2^p -complete for $\sigma \in \{\text{gr}, \text{co}, \text{st}, \text{pr}\}$;
- $c\text{-NV}_\sigma$ is Θ_2^p -complete for $\sigma \in \{\text{gr}, \text{co}, \text{st}, \text{pr}\}$;
- $s\text{-NV}_\sigma$ is Π_2^p -complete for $\sigma \in \{\text{gr}, \text{co}, \text{st}, \text{pr}\}$.

Notably, the verification problem for CKAF is computationally harder than that for *constrained iAF* (ciAF) (Mailly 2020) (and thus iAF) under all the considered (argumentation and subframework) semantics, except for $\delta\text{-PV}_{\text{pr}}$ (with $\delta \in \{c, s\}$) for which the verification problems for CKAF and ciAF belong to the same class (Σ_2^p). On the other side, it can be shown that the complexity of verification for KAF is as that of AF, as any KAF $\Omega = \langle A, R, \mathcal{T}, \mathcal{A} \rangle$ can be seen as a CKAF $\langle A, R, \mathcal{T}, \mathcal{A}, \Gamma = \emptyset \rangle$, whose set of constraints is empty, and $\langle \mathcal{T}, \mathcal{A}^* \rangle \in \sigma(\Omega)$ iff $\text{Arg}(\mathcal{A}^*) \in \sigma(\langle A, R \rangle)$.

5.3 Acceptance Problems

For AFs, the complexity of the credulous and skeptical acceptance problems has been investigated in (Dung 1995) for the grounded semantics, in (Dimopoulos and Torres 1996) for the stable semantics, and in (Dimopoulos and Torres 1996; Dunne and Bench-Capon 2002) for the preferred semantics. The complexity results for AFs are summarized in the second and third column of Table 2.

The following definition generalizes that of the credulous and skeptical acceptance problems for AF to the case of CKAF, analogously to what is done for (c)iAF.

Definition 12. Let $\Delta = \langle A, R, \mathcal{T}, \mathcal{A}, \Gamma \rangle$ be a CKAF, $\sigma \in \{\text{gr}, \text{co}, \text{pr}, \text{st}\}$ an argumentation semantics, $q(\vec{x})$ a UCQ, $\delta \in \{s, c\}$ a subframework semantics, and $\vec{a}$ a tuple of individuals. Then:

- the δ -possible credulous acceptance problem, denoted as $\delta\text{-PCA}_\sigma$, is the problem of deciding whether $\vec{a}$ is possibly credulously accepted;
- the δ -possible skeptical acceptance problem, denoted as $\delta\text{-PSA}_\sigma$, is the problem of checking whether $\vec{a}$ is possibly skeptically accepted;
- the δ -necessary credulous acceptance problem, denoted as $\delta\text{-NCA}_\sigma$, is the problem of checking whether $\vec{a}$ is necessary credulously accepted;
- the δ -necessary skeptical acceptance problem, denoted as $\delta\text{-NSA}_\sigma$, is the problem of checking whether $\vec{a}$ is necessary skeptically accepted.

The next theorem provides a tight characterization of the complexity for all the above-defined acceptance problems.

Theorem 3.

- $c\text{-PCA}_\sigma$ is Θ_2^p -complete for any $\sigma \in \{\text{gr}, \text{co}, \text{st}, \text{pr}\}$;
- $s\text{-PCA}_\sigma$ is Σ_2^p -complete for any $\sigma \in \{\text{gr}, \text{co}, \text{st}, \text{pr}\}$;
- $c\text{-NCA}_\sigma$ is i) Θ_2^p -complete for $\sigma = \text{gr}$, and ii) Π_2^p -complete for any $\sigma \in \{\text{co}, \text{st}, \text{pr}\}$;

σ	AF/KAF		iAF/ciAF				CKAF							
	CA_σ	SA_σ	PCA_σ	NCA_σ	PSA_σ	NSA_σ	$c-PCA_\sigma$	$s-PCA_\sigma$	$c-NCA_\sigma$	$s-NCA_\sigma$	$c-PSA_\sigma$	$s-PSA_\sigma$	$cNSA_\sigma$	$sNSA_\sigma$
gr	P	P	NP-c	coNP-c	NP-c	coNP-c	Θ_2^p -c	Σ_2^p -c	Θ_2^p -c	Π_2^p -c	Θ_2^p -c	Σ_2^p -c	Θ_2^p -c	Π_2^p -c
co	NP-c	P	NP-c	Π_2^p -c	NP-c	coNP-c	Θ_2^p -c	Σ_2^p -c	Π_2^p -c*	Π_2^p -c	Θ_2^p -c	Σ_2^p -c	Θ_2^p -c	Π_2^p -c
st	NP-c	coNP-c	NP-c	Π_2^p -c	Σ_2^p -c	coNP-c	Θ_2^p -c	Σ_2^p -c	Π_2^p -c*	Π_2^p -c	Σ_2^p -c*	Σ_2^p -c	Θ_2^p -c	Π_2^p -c
pr	NP-c	Π_2^p -c	NP-c	Π_2^p -c	Σ_3^p -c	Π_2^p -c	Θ_2^p -c	Σ_2^p -c	Π_2^p -c*	Π_2^p -c	Σ_3^p -c*	Σ_3^p -c*	Π_2^p -c*	Π_2^p -c

Table 2: Complexity of possible and necessary credulous/skeptical acceptance under $\sigma \in \{\text{gr}, \text{co}, \text{pr}, \text{st}\}$ for AF, KAF, iAF, ciAF, and CKAF. For any complexity class $\mathcal{C}$, $\mathcal{C}$ -c means $\mathcal{C}$ -complete. All results for CKAF (and KAF) are new. The results with asterisk (*) refer to cases where the complexity of CKAF is the same as that of (c)iAF; for the other cases, the complexity of CKAF is higher than that of (c)iAF.

- $s-NCA_\sigma$ is Π_2^p -complete for any $\sigma \in \{\text{gr}, \text{co}, \text{st}, \text{pr}\}$;
- $c-PSA_\sigma$ is i) Θ_2^p -complete for any $\sigma \in \{\text{gr}, \text{co}\}$, and
ii) Σ_2^p -complete for any $\sigma \in \{\text{st}, \text{pr}\}$;
- $s-PSA_\sigma$ is i) Σ_2^p -complete for any $\sigma \in \{\text{gr}, \text{co}, \text{st}\}$, and
ii) Σ_3^p -complete for $\sigma = \text{pr}$;
- $c-NSA_\sigma$ is i) Θ_2^p -complete for any $\sigma \in \{\text{gr}, \text{co}, \text{st}\}$, and
ii) Π_2^p -complete for $\sigma = \text{pr}$.
- $s-NSA_\sigma$ is Π_2^p -complete for any $\sigma \in \{\text{gr}, \text{co}, \text{st}, \text{pr}\}$.

In brief, except for the 7 cases highlighted with an asterisk (*) in Table 2 for which the complexity of CKAF is the same as that of (c)iAF, for all the other 25 cases considered the complexity of CKAF is higher than that of (c)iAF, even under cardinality semantics whose complexity is systematically lower than or equal to that of the set semantics.

6 CKAF versus iAF-based Frameworks

As observed earlier, there are some connections between CKAF and iAF-based Frameworks (Baumeister et al. 2021; Fazzinga, Flesca, and Furfaro 2020; Mailly 2020). Acceptance problems are defined similarly, in the sense that the semantics of iAF-based frameworks is defined by the set of (valid) completions, whereas that of CKAF is defined by the set of subframeworks (that can be viewed as a subset of completions by assuming arguments/attacks of CKAF to be uncertain). We investigate the relationship between CKAF and *constrained iAF* (ciAF) (Mailly 2020), showing that ciAF, as well as its restrictions iAF and *correlated iAF* (Fazzinga, Flesca, and Furfaro 2020), are special cases of CKAF.⁴

6.1 Constrained iAF

A *constrained incomplete (abstract) Argumentation Framework* (ciAF) is a tuple $\Psi = \langle A, B, R, T, C \rangle$, where A and B are disjoint sets of arguments, and R and T are disjoint sets of attacks between arguments in $A \cup B$. Arguments in A and attacks in R are said to be *certain*, while arguments in B and attacks in T are said to be *uncertain* (Baumeister




Figure 2: ciAF of Example 11 (dashed elements are uncertain).

et al. 2018). Moreover, $C \in \mathcal{L}_{A \cup B}$ is a propositional formula (called constraint), defined over the propositional language $\mathcal{L}_A$ built over the propositional atoms in $\{arg(a) \mid a \in A\} \cup \{att(a, b) \mid a, b \in A\}$. A ciAF compactly represents alternative AF scenarios, called *valid completions*.

Definition 13 (Valid Completions). A *valid completion* for a ciAF $\Psi = \langle A, B, R, T, C \rangle$ is an AF $\Lambda = \langle A', R' \rangle$ where $A \subseteq A' \subseteq A \cup B$, $R \cap (A' \times A') \subseteq R' \subseteq (R \cup T) \cap (A' \times A')$, and Λ satisfies C . The set of valid completions of Ψ is denoted by $valid-comp(\Psi)$.

Verification and acceptance problems, originally introduced in (Baumeister, Neugebauer, and Rothe 2018; Baumeister et al. 2021) for iAF, are defined analogously to the ones of CKAF, with the only difference that instead of considering subframeworks and general UCQs, valid completions and single goal arguments are taken into account.

Example 11. Consider the ciAF $\Psi = \langle A = \{b, c, d\}, B = \{a\}, R = \{(a, c), (b, c), (c, d)\}, T = \{(a, b)\}, C = arg(a) \Rightarrow att(a, b) \rangle$ shown in Figure 2, where arguments b, c, d and attacks $(a, c), (b, c)$ and (c, d) are certain, whereas argument a and attack (a, b) are uncertain.

Ψ has 2 valid completions, that are:

- $\Lambda_1 = \langle \{a, b, c, d\}, \{(a, b), (a, c), (b, c), (c, d)\} \rangle$; and
- $\Lambda_2 = \langle \{b, c, d\}, \{(b, c), (c, d)\} \rangle$.

For $\sigma \in \{\text{gr}, \text{co}, \text{st}, \text{pr}\}$, Λ_1 (resp. Λ_2) has only one extension $E_1 = \{a, d\}$ (resp. $E_2 = \{b, d\}$). Thus, a, b, d are possibly credulously/skeptically accepted, while only d is necessarily credulously/skeptically accepted. $\square$

6.2 Relationship between CKAF and ciAF

We start by showing that every ciAF can be rewritten into an extensions-equivalent CKAF, modulo meta-elements. The equivalence between a ciAF $\Psi = \langle A, B, R, T, C \rangle$ and a CKAF $\Delta = \langle A', R', \mathcal{T}, \mathcal{A}, \Gamma \rangle$ derived from Ψ is in the sense that (i) $A \cup B \subseteq A'$ and $R \cup T \subseteq R'$ (i.e., arguments and attacks in Ψ also occur in Δ , but Δ may contain additional meta-arguments and meta-attacks), and (ii) for every valid completion $\Lambda \in valid-comp(\Psi)$ there exists an s-subframework $\langle A'', R'', \mathcal{T}, \mathcal{A}'' \rangle \in \mathcal{S}^s(\Delta)$ (resp. c-subframework $\langle A'', R'', \mathcal{T}, \mathcal{A}'' \rangle \in \mathcal{S}^c(\Delta)$) such that $\sigma(\Lambda) =$

⁴While correlated iAF considers different fragments of propositional logic (PL), with the aim of studying complexity and expressivity of those fragments, constrained iAF considers general PL formulae. As PL formulae can be rewritten using minimal sets of operators (e.g., $\wedge$ and $\neg$), the two proposals are equivalent in the sense that, although constrained iAF is more general, there are fragments of correlated iAF that have the same expressivity of general PL. Finally, an iAF is a special ciAF without constraints.

$\{E \cap (A \cup B) \mid E \in \sigma(\langle A'', R'' \rangle)\}$, and vice versa. Such equivalence under the set subframework (resp. cardinality subframework) semantics between a ciAF Ψ and a (derived) CKAF Δ is denoted by $\Psi \approx_s \Delta$ (resp. $\Psi \approx_c \Delta$).

Theorem 4. *For any ciAF $\Psi = \langle A, B, R, T, C \rangle$, it holds that $\Psi \approx_s \Delta_\Psi$, where $\Delta_\Psi = \langle A_\Psi, R \cup T, \emptyset, af(A_\Psi, R_\Psi), \text{true} \Rightarrow c_1 \wedge c_2 \wedge c_3 \wedge C \rangle$ is the CKAF obtained from Δ as follows: $A_\Psi = A \cup B \cup \bar{B} \cup \bar{R}$, where $\bar{B} = \{\bar{b} \mid b \in B\}$; $\bar{R} = \{\bar{ab} \mid (a, b) \in R \wedge (a \in B \vee b \in B)\} \cup \{\bar{ab} \mid (a, b) \in T\}$; and*

$$\begin{aligned} c_1 : & \bigwedge_{a \in A} \text{arg}(a) \wedge \bigwedge_{(a,b) \in (R \cap (A \times A))} \text{att}(a, b) \\ c_2 : & \frac{\bigwedge_{a \in B} (\text{arg}(a) \oplus \text{arg}(\bar{a})) \wedge \bigwedge_{(a,b) \in T} \text{att}(a, b) \oplus \text{arg}(\bar{ab})}{\bigwedge_{(a,b) \in R \wedge (\{a,b\} \cap B \neq \emptyset)} \text{att}(a, b) \oplus \text{arg}(\bar{ab})} \\ c_3 : & \frac{\bigwedge_{(a,b) \in (R \cap (B \times A))} (\text{arg}(a) \equiv \text{att}(a, b)) \wedge \bigwedge_{(a,b) \in (R \cap (A \times B))} (\text{arg}(b) \equiv \text{att}(a, b)) \wedge \bigwedge_{(a,b) \in (R \cap (B \times B))} (\text{arg}(a) \wedge \text{arg}(b) \equiv \text{att}(a, b)) \wedge \bigwedge_{(a,b) \in T} (\text{att}(a, b) \implies \text{arg}(a) \wedge \text{arg}(b))}{\end{aligned}$$

Example 12. Consider the ciAF $\Psi = \langle A, B, R, T, C \rangle$ of Example 11. The CKAF derived from Ψ is $\Delta_\Psi = \langle A_\Psi = A \cup B \cup \{\bar{a}, \bar{ab}, \bar{ac}\}, R_\Psi = R \cup T, \emptyset, af(A_\Psi, R_\Psi), \Gamma \rangle$ where:

$$\begin{aligned} \Gamma &= \text{true} \Rightarrow c_1 \wedge c_2 \wedge c_3 \wedge C, \text{ with} \\ c_1 &= \text{arg}(b) \wedge \text{arg}(c) \wedge \text{arg}(d) \wedge \text{att}(b, c) \wedge \text{att}(c, d); \\ c_2 &= (\text{arg}(a) \oplus \text{arg}(\bar{a})) \wedge (\text{att}(a, b) \oplus \text{arg}(\bar{ab})) \\ &\quad \wedge (\text{att}(a, c) \oplus \text{arg}(\bar{ac})); \\ c_3 &= \text{arg}(a) \equiv \text{att}(a, c); \end{aligned}$$

Intuitively, (i) c_1 states that certain arguments and certain attacks, whose end vertices are certain, must belong to all subframeworks; (ii) c_2 states that for the uncertain argument a (resp. uncertain attack (a, b) and (a, c)) exactly one between a and $\bar{a}$ (resp. between $\text{att}(a, b)$ and $\text{arg}(\bar{ab})$, $\text{att}(a, c)$ and $\text{arg}(\bar{ac})$) appears in any subframework (in the formula, $\oplus$ denotes the XOR operator); (iii) c_3 states that the certain attack (a, c) must belong to a subframework iff a (which is uncertain in the iAF) belong to it.

Δ_Ψ has two s-subframeworks (also c-subframeworks) $\Omega^i = \langle A^i, R^i, \emptyset, af(A^i, R^i) \rangle$ (with $i \in \{1, 2\}$), where:

- $A^1 = \{a, b, c, d\}$, and $R^1 = \{(a, b), (b, c), (a, c), (c, d)\}$;
- $A^2 = \{\bar{a}, b, c, d, \bar{ab}, \bar{ac}\}$, and $R^2 = \{(b, c), (c, d)\}$,

which one-to-one correspond to the valid completions of ciAF Ψ , modulo meta-elements. Also, the set of σ -extensions (with $\sigma \in \{\text{gr}, \text{co}, \text{st}, \text{pr}\}$) of Ω^i is the same as that of Δ_i (of Example 11), modulo the meta-arguments $\bar{a}$, $\bar{ab}$, and $\bar{ac}$, introduced in the rewriting, for $i \in \{1, 2\}$. $\square$

The result of Theorem 4 can be viewed in the context of *realizability* in that realizability has also been used to compare the expressiveness of two frameworks (Linsbichler, Pührer, and Strass 2016; Dunne et al. 2015; Baumann et al. 2014; Strass 2015; Pührer 2020). Following the approach of (Gogic et al. 1995), where it is stated that a framework F is at least as expressive as framework G if and only if every knowledge base in G has an equivalent knowledge base in F (obtainable from G in polynomial time), Theorem 4 states that CKAF is at least as expressive as ciAF. Under the realizability viewpoint, our result entails that the set of extensions of a given ciAF can be modelled by means of a CKAF

(obtainable from the ciAF in polynomial time), regardless of the specific semantics. Moreover, it is worth noting that, using the terminology introduced in (Brewka, Dunne, and Woltran 2011), we can say that ciAF can be *simulated* in CKAF. However, it is important to note that the inverse does not hold, that is, CKAF is strictly more expressive than ciAF. In fact, as also backed by our complexity results, several CKAF reasoning problems (e.g., credulous/skeptical acceptance) cannot be reduced to their counterparts in ciAF.

7 Related Work

The concept of CKAF subframework shares some basic ideas with that of KB repairs (Arenas, Bertossi, and Chomicki 1999; Bertossi 2011; Calautti et al. 2022a; Calautti et al. 2022b). However, here we aim to select (maximal) subsets of the underlying AF that represent feasible solutions consistently with the additional knowledge, rather than identifying and rectifying inconsistencies in the data. Computing repairs for AFs which are inconsistent, in the sense that they do not have any accepted argument, has been investigated in (Ulbricht and Baumann 2019), where restoring consistency is achieved via dropping a minimal set of arguments or attacks. A similar approach has been defined for Assumption-Based Argumentation Frameworks in (Rapberger and Ulbricht 2024). In (Takahashi and Miwa 2023) it is studied how to make an AF with no stable extensions into one having it by adding a new argument (called ‘repair’).

The relationship between argumentation and repairs in inconsistent databases (DBs) has been investigated in (Mahmood, Hecher, and Ngomo 2025) with the aim to show that classical AF semantics can be reduced to the computation of repairs over a DB inconsistent w.r.t. a set of functional and inclusion dependencies. Moreover, in (Bienvenu and Bourgaux 2020), it has been shown that AF stable extensions correspond to Pareto-optimal repairs. In contrast, our aim is to first enrich AF with a KB, and then to provide means for selecting subframeworks that are of interest to the user.

Further related approaches address the so-called *enforcing* and *minimal change* problem dealing with the question of whether it is possible to add (minimal) new information so that a desired set of arguments becomes an extension or at least a subset of an extension (Baumann 2012; Kim, Ordyniak, and Szeider 2013; Wallner, Niskanen, and Järvisalo 2017). These works are all based on the addition of new information. We do not allow modifications that result in the addition of arguments and attacks, but we are interested in (minimal) removals ensuring satisfaction of user-defined constraints. The impact of adding or removing an argument on the set of extensions has been studied in (Cayrol, de Saint-Cyr, and Lagasquie-Schiex 2010).

Extending AF by associating a claim to each argument representing its conclusion has been investigated in (Dvorák, Rapberger, and Woltran 2020; Dvorák and Woltran 2020; Dvorák et al. 2023) where claim-augmented AFs are introduced. A claim-augmented AF can be modeled by a KAF whose TBox is empty and whose ABox contains assertions associating arguments to claims. Two approaches have been considered for defining claim-augmented AF semantics.

The *inherited* semantics considers the extensions of the underlying AF before interpreting arguments by their claims, while the *claim-level* semantics redefines argumentation semantics in terms of the arguments' claims. We follow the inherited approach in defining KAF semantics, as often done in structured argumentation (Modgil and Prakken 2014; Bondarenko, Toni, and Kowalski 1993). As a result, KAF generalizes claim-augmented AF under *inherited* semantics, as backed by our complexity analysis.

Interestingly, also Preference-based AF (PAF) under the preference reduction approach (Amgoud and Cayrol 2002; Amgoud and Vesic 2014; Kaci et al. 2021; Bernreiter et al. 2024; Bernreiter, Dvorák, and Woltran 2024) can be reduced to CKAF where PAF's *critical* attacks are tackled similarly to VAF's *unsuccessful* attacks (Dunne and Bench-Capon 2004), as shown in what follows. Any VAF $\langle A, R, V, \text{VAL}, \text{VALPREF} \rangle$ can be seen as a CKAF $\langle A, R, \emptyset, af(A, R) \cup \mathcal{F}, \Gamma \rangle$ where $\mathcal{F} = \{ \text{val}(x, \text{VAL}(x)) \mid x \in A \} \cup \{ \text{pref}(x, y) \mid \{x, y\} \subseteq V \wedge \text{VALPREF}(x, y) \}$, and $\Gamma = \forall x, y, u, v. \text{att}(x, y) \wedge \text{val}(x, u) \wedge \text{val}(y, v) \wedge \text{pref}(v, u) \Rightarrow \text{false}$. Here, the unique δ -subframework corresponds to an AF without unsuccessful attacks, emulating VAF semantics.

While Dung's framework is mainly concerned with acceptability of arguments, structured argumentation (Besnard et al. 2014) includes an argument generation process or instantiation, upon which semantics can be deployed to find acceptable arguments or conclusions thereof. In particular, ASPIC⁺ (Modgil and Prakken 2014) and ABA (Bondarenko, Toni, and Kowalski 1993; Toni 2014) build the set of all possible arguments from a KB and then rely on using one of the possible Dung semantics to decide on the acceptance of arguments, while Logic-Based Deductive Argumentation (Besnard and Hunter 2014) and DeLP (García and Simari 2014) focus on building the arguments involved in answering the query. Differently from structured argumentation approaches, where the internal logical structure of arguments determines the relationship between them, KAF and CKAF build on Dung's framework, where the relationships between arguments are explicit (not inferred from a KB), and enrich it by DLs that handle structured domain knowledge, leveraging AF semantics for management of conflicts.

8 Conclusion

We have introduced a new framework combining abstract argumentation with description logics. Although for the sake of presentation we focused on *DL-Lite_R*, it can be shown that our results hold for all query rewritable *DL-Lite* fragments (Calvanese et al. 2007; Bienvenu 2016), as they hold as long as the query (and the constraints) can be evaluated in polynomial time. Future work should be devoted to the investigation of integrating AF with more general DL knowledge bases, by considering both more powerful query languages (Bienvenu, Manière, and Thomazo 2020; Bienvenu, Manière, and Thomazo 2021) and more powerful DL languages for axioms' definition (Cima, Lenzerini, and Poggi 2020; Cima et al. 2025). The combination of the two frameworks could develop interesting researches as it has the potential to develop reasoning tasks not supported by single frameworks such as: (i) Defeasible reason-

ing over ontologies, enabling systems to withdraw conclusions when new, contradictory evidence arises; (ii) Explainable AI, where arguments built over DL-based knowledge can provide structured justifications for conclusions; (iii) Interoperability in knowledge integration, managing inconsistencies when merging heterogeneous knowledge bases.

Although this paper primarily focuses on formal foundations, bridging this theoretical framework with practical implementations is important. While extending current AF solvers to compute KAF semantics is not excluded from our complexity results, the same does not hold for CKAF whose underlying problems are strictly harder than AF. A possible solution would be that of hybrid solving approaches that combine existing DL reasoners (e.g., Pellet, HermiT) for KB reasoning and AF solvers (e.g., those of the ICCMA competition) for extensions computation. We believe that modular integration of these tools—possibly via interleaved or layered reasoning—could offer an effective path forward. This is a direction that we plan to explore in future work.

Acknowledgements

This work is dedicated to Professor Maurizio Lenzerini on the occasion of his 70th birthday, in recognition of his seminal contributions to Knowledge Representation, particularly the introduction of the DL-Lite family of description logics.

We acknowledge the support from project Tech4You (ECS0000009), and PNRR MUR projects FAIR (PE0000013) and SERICS (PE0000014).

References

- Abiteboul, S.; Hull, R.; and Vianu, V. 1995. *Foundations of Databases*. Addison-Wesley.
- Alfano, G.; Greco, S.; Parisi, F.; and Trubitsyna, I. 2023. Abstract argumentation framework with conditional preferences. In *Proc. of AAAI*, 6218–6227.
- Alfano, G.; Greco, S.; Mandaglio, D.; Parisi, F.; and Trubitsyna, I. 2024a. Abstract argumentation frameworks with strong and weak constraints. *Artif. Intell.* 336:104205.
- Alfano, G.; Greco, S.; Parisi, F.; and Trubitsyna, I. 2024b. Cyclic supports in recursive bipolar argumentation frameworks: Semantics and LP mapping. *Theory Pract. Log. Program.* 24(4):921–941.
- Amgoud, L., and Cayrol, C. 1998. On the acceptability of arguments in preference-based argumentation. In *Proc. of UAI*, 1–7.
- Amgoud, L., and Cayrol, C. 2002. Inferring from inconsistency in preference-based argumentation frameworks. *J. Autom. Reason.* 29(2):125–169.
- Amgoud, L., and Vesic, S. 2014. Rich preference-based argumentation frameworks. *Int. J. Approx. Reason.* 55(2):585–606.
- Amgoud, L.; Dimopoulos, Y.; and Moraitis, P. 2007. A unified and general framework for argumentation-based negotiation. In *Proc. of AAMAS*, 158.
- Arenas, M.; Bertossi, L. E.; and Chomicki, J. 1999. Consistent query answers in inconsistent databases. In *Proc. of PODS*, 68–79.

- Arieli, O. 2015. Conflict-free and conflict-tolerant semantics for constrained argumentation frameworks. *J. Appl. Log.* 13(4):582–604.
- Artale, A.; Calvanese, D.; Kontchakov, R.; and Zakharyashev, M. 2009. The dl-lite family and relations. *J. Artif. Intell. Res.* 36:1–69.
- Baumann, R.; Dvorák, W.; Linsbichler, T.; Strass, H.; and Woltran, S. 2014. Compact argumentation frameworks. In *Proc. of ECAI*, volume 263, 69–74.
- Baumann, R. 2012. What does it take to enforce an argument? minimal change in abstract argumentation. In *Proc. of ECAI*, volume 242, 127–132.
- Baumeister, D.; Neugebauer, D.; Rothe, J.; and Schadrack, H. 2018. Verification in incomplete argumentation frameworks. *Artif. Intell.* 264:1–26.
- Baumeister, D.; Järvisalo, M.; Neugebauer, D.; Niskanen, A.; and Rothe, J. 2021. Acceptance in incomplete argumentation frameworks. *Artif. Intell.* 103470.
- Baumeister, D.; Neugebauer, D.; and Rothe, J. 2018. Credulous and skeptical acceptance in incomplete argumentation frameworks. In *Proc. of COMMA*, 181–192.
- Bernreiter, M.; Dvorák, W.; Rapberger, A.; and Woltran, S. 2024. The effect of preferences in abstract argumentation under a claim-centric view. *J. Artif. Intell. Res.* 81:203–262.
- Bernreiter, M.; Dvorák, W.; and Woltran, S. 2024. Abstract argumentation with conditional preferences. *Argument Comput.* 15(2):161–189.
- Bertossi, L. E. 2011. *Database Repairing and Consistent Query Answering*. Synthesis Lectures on Data Management. Morgan & Claypool Publishers.
- Besnard, P., and Hunter, A. 2014. Constructing argument graphs with deductive arguments: A tutorial. *Argument & Computation* 5(1):5–30.
- Besnard, P.; Garcia, A. J.; Hunter, A.; Modgil, S.; Prakken, H.; Simari, G. R.; and Toni, F. 2014. Introduction to structured argumentation. *Argument & Computation – Special Issue: Tutorials on Structured Argumentation* 5(1):1–4.
- Bienvenu, M., and Bourgaux, C. 2020. Querying and repairing inconsistent prioritized knowledge bases: Complexity analysis and links with abstract argumentation. In *Proc. of KR*, 141–151.
- Bienvenu, M., and Ortiz, M. 2015. Ontology-mediated query answering with data-tractable description logics. In *Proc. of RW*, volume 9203 of *Lecture Notes in Computer Science*, 218–307. Springer.
- Bienvenu, M.; Manière, Q.; and Thomazo, M. 2020. Answering counting queries over dl-lite ontologies. In *Proc. of IJCAI*, 1608–1614.
- Bienvenu, M.; Manière, Q.; and Thomazo, M. 2021. Cardinality queries over dl-lite ontologies. In *Proc. of IJCAI*, 1801–1807.
- Bienvenu, M. 2016. Ontology-mediated query answering: Harnessing knowledge to get more from data. In *Proc. of IJCAI*, 4058–4061.
- Bondarenko, A.; Toni, F.; and Kowalski, R. A. 1993. An assumption-based framework for non-monotonic reasoning. In *Proc. of LPNMR*, 171–189.
- Brachman, R. J., and Levesque, H. J. 2004. *Knowledge Representation and Reasoning*. Elsevier.
- Brewka, G., and Woltran, S. 2010. Abstract dialectical frameworks. In *Proc. of KR*.
- Brewka, G.; Strass, H.; Ellmauthaler, S.; Wallner, J. P.; and Woltran, S. 2013. Abstract dialectical frameworks revisited. In *Proc. of IJCAI*, 803–809.
- Brewka, G.; Dunne, P. E.; and Woltran, S. 2011. Relating the semantics of abstract dialectical frameworks and standard afs. In *Proc. of IJCAI*, 780–785.
- Calautti, M.; Greco, S.; Molinaro, C.; and Trubitsyna, I. 2022a. Preference-based inconsistency-tolerant query answering under existential rules. *Artif. Intell.* 312:103772.
- Calautti, M.; Greco, S.; Molinaro, C.; and Trubitsyna, I. 2022b. Query answering over inconsistent knowledge bases: A probabilistic approach. *Theor. Comput. Sci.* 935:144–173.
- Calvanese, D.; Giacomo, G. D.; Lembo, D.; Lenzerini, M.; and Rosati, R. 2007. Tractable reasoning and efficient query answering in description logics: The *DL-Lite* family. *J. Autom. Reason.* 39(3):385–429.
- Caminada, M. 2006. Semi-stable semantics. In *Proc. of COMMA*, 121–130.
- Cayrol, C.; Fandinno, J.; del Cerro, L. F.; and Lagasquie-Schiex, M. 2018. Structure-based semantics of argumentation frameworks with higher-order attacks and supports. In *Proc. of COMMA*, 29–36.
- Cayrol, C.; de Saint-Cyr, F. D.; and Lagasquie-Schiex, M. 2010. Change in abstract argumentation frameworks: Adding an argument. *J. Artif. Intell. Res.* 38:49–84.
- Cima, G.; Console, M.; Delfino, R. M.; Lenzerini, M.; and Poggi, A. 2025. Answering conjunctive queries with safe negation and inequalities over RDFS knowledge bases. In *Proc. of AAI*, 14824–14831.
- Cima, G.; Lenzerini, M.; and Poggi, A. 2020. Answering conjunctive queries with inequalities in *DL-Lite*. In *Proc. of AAI*, 2782–2789.
- Cohen, A.; Gottifredi, S.; Garcia, A. J.; and Simari, G. R. 2015. An approach to abstract argumentation with recursive attack and support. *J. Appl. Log.* 13(4):509–533.
- Coste-Marquis, S.; Devred, C.; and Marquis, P. 2006. Constrained argumentation frameworks. In *Proc. of KR*, 112–122.
- Dimopoulos, Y., and Torres, A. 1996. Graph theoretical structures in logic programs and default theories. *Theor. Comput. Sci.* 170(1-2):209–244.
- Dimopoulos, Y.; Mailly, J.; and Moraitis, P. 2019. Argumentation-based negotiation with incomplete opponent profiles. In *Proc. of AAMAS*, 1252–1260.
- Dung, P. M., and Thang, P. M. 2010. Towards (probabilistic) argumentation for jury-based dispute resolution. In *Proc. of COMMA*, 171–182.

- Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artif. Intell.* 77:321–358.
- Dunne, P. E., and Bench-Capon, T. J. M. 2002. Coherence in finite argument systems. *Artif. Intell.* 141(1/2):187–203.
- Dunne, P. E., and Bench-Capon, T. J. M. 2004. Complexity in value-based argument systems. In *Proc. of JELIA*, 360–371.
- Dunne, P. E.; Dvorák, W.; Linsbichler, T.; and Woltran, S. 2015. Characteristics of multiple viewpoints in abstract argumentation. *Artif. Intell.* 228:153–178.
- Dvorak, W., and Dunne, P. E. 2017. Computational problems in formal argumentation and their complexity. *FLAP* 4(8).
- Dvorák, W., and Woltran, S. 2020. Complexity of abstract argumentation under a claim-centric view. *Artif. Intell.* 285:103290.
- Dvorák, W.; Greßler, A.; Rapberger, A.; and Woltran, S. 2023. The complexity landscape of claim-augmented argumentation frameworks. *Artif. Intell.* 317:103873.
- Dvorák, W.; Rapberger, A.; and Woltran, S. 2020. Argumentation semantics under a claim-centric view: Properties, expressiveness and relation to setafs. In Calvanese, D.; Erdem, E.; and Thielscher, M., eds., *Proc. of KR*, 341–350.
- Fazzinga, B.; Flesca, S.; and Furfaro, F. 2020. Revisiting the notion of extension over incomplete abstract argumentation frameworks. In *Proc. of IJCAI*, 1712–1718.
- Fazzinga, B.; Flesca, S.; and Furfaro, F. 2021. Reasoning over argument-incomplete aafs in the presence of correlations. In *Proc. of IJCAI*, 189–195.
- Gabbay, D.; Giacomini, M.; Simari, G. R.; and Thimm, M., eds. 2021. *Handbook of Formal Argumentation*, volume 2. College Publications.
- Gabbay, D.; Kern-Isberner, G.; Simari, G. R.; and Thimm, M., eds. 2024. *Handbook of Formal Argumentation*, volume 3. College Publications.
- García, A. J., and Simari, G. R. 2014. Defeasible logic programming: DeLP-servers, contextual queries, and explanations for answers. *Argument & Computation* 5(1):63–88.
- Gogic, G.; Kautz, H. A.; Papadimitriou, C. H.; and Selman, B. 1995. The comparative linguistics of knowledge representation. In *Proc. of IJCAI*, 862–869.
- Gottifredi, S.; Cohen, A.; Garcia, A. J.; and Simari, G. R. 2018. Characterizing acceptability semantics of argumentation frameworks with recursive attack and support relations. *Artif. Intell.* 262:336–368.
- Hunter, A. 2012. Some foundations for probabilistic abstract argumentation. In *Proc. of COMMA*, 117–128.
- Kaci, S.; van der Torre, L. W. N.; Vesic, S.; and Villata, S. 2021. Preference in abstract argumentation. In *Handbook of Formal Argumentation*, volume 2. College Publications. chapter 3.
- Kim, E. J.; Ordyniak, S.; and Szeider, S. 2013. The complexity of repairing, adjusting, and aggregating of extensions in abstract argumentation. In *Proc. of TAFE*, volume 8306, 158–175.
- Li, H.; Oren, N.; and Norman, T. J. 2011. Probabilistic argumentation frameworks. In *Proc. of TAFE*, 1–16.
- Linsbichler, T.; Pührer, J.; and Strass, H. 2016. A uniform account of realizability in abstract argumentation. In *Proc. of ECAI*, volume 285 of *Frontiers in Artificial Intelligence and Applications*, 252–260.
- Mahmood, Y.; Hecher, M.; and Ngomo, A. N. 2025. Dung’s argumentation framework: Unveiling the expressive power with inconsistent databases. In *Proc. of AFAI*, 15058–15066.
- Mailly, J. 2020. Possible controllability of control argumentation frameworks. In *Proc. of COMMA*, 283–294.
- Mailly, J. 2024. Constrained incomplete argumentation frameworks: Expressiveness, complexity and enforcement. *AI Commun.* 37(3):299–322.
- Modgil, S., and Prakken, H. 2013. A general account of argumentation with preferences. *Artif. Intell.* 195:361–397.
- Modgil, S., and Prakken, H. 2014. The *ASPIC⁺* framework for structured argumentation: A tutorial. *Argument & Computation* 5(1):31–62.
- Nouioua, F., and Risch, V. 2011. Argumentation frameworks with necessities. In *Proc. of SUM*.
- Nouioua, F. 2013. Afs with necessities: Further semantics and labelling characterization. In *Proc. of SUM*, 120–133.
- Prakken, H. 2009. Models of persuasion dialogue. In *Argumentation in Artificial Intelligence*. Springer. 281–300.
- Pührer, J. 2020. Realizability of three-valued semantics for abstract dialectical frameworks. *Artif. Intell.* 278.
- Rapberger, A., and Ulbricht, M. 2024. Repairing assumption-based argumentation frameworks. In *Proc. of KR*.
- Strass, H. 2015. Expressiveness of two-valued semantics for abstract dialectical frameworks. *J. Artif. Intell. Res.* 54:193–231.
- Takahashi, K., and Miwa, H. 2023. Topological conditions and solutions for repairing argumentation frameworks. In *Proc. of CLAR*, volume 14156 of *Lecture Notes in Computer Science*, 101–118. Springer.
- Toni, F. 2014. A tutorial on assumption-based argumentation. *Argument & Computation* 5(1):89–117.
- Ulbricht, M., and Baumann, R. 2019. If nothing is accepted - repairing argumentation frameworks. *J. Artif. Intell. Res.* 66:1099–1145.
- Villata, S.; Boella, G.; Gabbay, D. M.; and van der Torre, L. W. N. 2012. Modelling defeasible and prioritized support in bipolar argumentation. *Ann. Math. Artif. Intell.* 66(1-4).
- W3C. 2012. OWL 2 Web Ontology Language Profiles (second edition). <https://www.w3.org/TR/owl2-profiles/>.
- Wallner, J. P.; Niskanen, A.; and Järvisalo, M. 2017. Complexity results and algorithms for extension enforcement in abstract argumentation. *J. Artif. Intell. Res.* 60:1–40.