Epistemic Abstract Argumentation Framework: Formal Foundations, Computation and Complexity

Gianvincenzo Alfano
University of Calabria
Rende, Italy
g.alfano@dimes.unical.it

Sergio Greco
University of Calabria
Rende, Italy
greco@dimes.unical.it

Francesco Parisi
University of Calabria
Rende, Italy
fparisi@dimes.unical.it

Irina Trubitsyna
University of Calabria
Rende, Italy
i.trubitsyna@dimes.unical.it

ABSTRACT

Dung's Abstract Argumentation Framework (AAF) has emerged as a central formalism in AI for modeling disputes among agents. In this paper, we introduce an extension of Dung's framework, called Epistemic Abstract Argumentation Framework (EAAF), which enhances AAF by allowing the representation of some pieces of epistemic knowledge. We generalize the concept of attack in AAF, introducing *strong* and *weak epistemic attacks* in EAAF, whose intuitive meaning is that an attacked argument is epistemically accepted only if the attacking argument is possibly or certainly rejected, respectively. We provide an intuitive semantics for EAAF that naturally extends that for AAF, and give an algorithm that enables the computation of epistemic extensions by using AAF-solvers. Finally, we analyze the complexity of the following argumentation problems: *verification*, i.e. checking whether a set of arguments is an epistemic extension; *existence*, i.e. checking whether there is at least one (non-empty) epistemic extension; and *acceptance*, i.e. checking whether an argument is epistemically accepted, under well-known argumentation semantics (i.e. grounded, complete, and preferred).

KEYWORDS

Formal Argumentation; Epistemic Argumentation; Complexity.

ACM Reference Format:
Gianvincenzo Alfano, Sergio Greco, Francesco Parisi, and Irina Trubitsyna. 2023. Epistemic Abstract Argumentation Framework: Formal Foundations, Computation and Complexity. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 9 pages.

1 INTRODUCTION

In the last decades, Argumentation [17, 22, 57] has become an important research field in the area of autonomous agents and multi-agent systems [56]. Argumentation has applications in several contexts, including modeling dialogues, negotiation [13, 32], and persuasion [52]. It has been widely used to model agents' interactions [14, 28, 49, 51], especially in the context of debates [36, 45, 53].

Dung's Abstract Argumentation Framework (AAF) is a simple yet powerful formalism for modeling disputes between two or more agents [33]. An AAF consists of a set of *arguments* and a binary *attack* relation over the set of arguments that specifies the interactions between arguments: intuitively, if argument a attacks argument b, then b is acceptable only if a is not. Hence, arguments are abstract entities whose status is entirely determined by the attack relation. An AAF can be seen as a directed graph, whose nodes represent arguments and edges represent attacks. Several argumentation semantics—e.g. *grounded* ($\mathtt{gr}$), *complete* ($\mathtt{co}$), *preferred* ($\mathtt{pr}$), and *stable* ($\mathtt{st}$) [33]—have been defined for AAF, leading to the characterization of *σ-extensions*, that intuitively consist of the sets of arguments that can be collectively accepted under semantics $\sigma \in \{\mathtt{gr}, \mathtt{co}, \mathtt{pr}, \mathtt{st}\}$.

Example 1. Consider an AAF Λ=$\langle\{\mathsf{a}, \mathsf{b}\}, \{(\mathsf{a}, \mathsf{b}), (\mathsf{b}, \mathsf{a})\}\rangle$ whose corresponding graph is shown in Figure 1(left). Λ describes the following scenario. A party planner invites Alice (a) and Bob (b) to join a party. Due to their old rivalry (*i*) Alice replies that she will not join the party if Bob does, and (*ii*) Bob replies that he will not join the party if Alice does. This situation can be modeled by AAF Λ, where an argument x states that "*(the person whose initial is)* x *joins the party*". Under the preferred semantics, there are two extensions $E_1 = \{\mathsf{a}\}$ and $E_2 = \{\mathsf{b}\}$ stating that only Alice or only Bob will attend the party, respectively. □

Thus, as prescribed by E_1 and E_2, in the previous example we have that the participation of Alice and Bob to the party is uncertain. To deal with uncertain information represented by the presence of multiple extensions, credulous and skeptical reasoning has been introduced. Specifically, an argument is credulously true (or accepted) if there exists an extension containing the argument, whereas an argument is skeptically true if it occurs in all extensions. However, uncertain information in AAF under multiple-status semantics proposed so far cannot be exploited to determine the status of arguments (which in turn influences the status of other arguments) by taking into account the information given by the whole set of extensions, as in the case of credulous and skeptical acceptance. To overcome such a situation, and thus provide a natural and compact way for expressing such kind of conditions, in this paper we propose the use of epistemic arguments and attacks. Informally, epistemic attacks allow considering all extensions and not only the current one. Thus, an epistemic attack from a to b is such that a defeats b if a occurs in at least one extension (*strong epistemic attack*) or in all extensions and at least one (*weak epistemic attack*).

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

Figure 1: AAF Λ of Example 1 (left) and EAAF Δ of Example 2 (right).

Example 2. Consider the AAF Λ of Example 1 and assume there are two more people: Carol ($\mathtt{c}$) and David ($\mathtt{d}$). Carol's answer is that she will not attend the party if it is sure (i.e. it is skeptically true) that Alice will, whereas David answers that he will not attend the party if the participation of Bob is possible (i.e. it is credulously true). Intuitively, the party planner should conclude that, as the participation of both Alice and Bob is uncertain, Carol will attend the party, whereas David will not.

This situation can be modeled by means of the Epistemic AAF (EAAF) shown in Figure 1(right) where $\mathtt{a}$ defeats $\mathtt{c}$ with a weak epistemic attack, whereas $\mathtt{b}$ defeats $\mathtt{d}$ with a strong epistemic attack (we use the two kinds of edges represented in the figure to denote weak and strong epistemic attacks). Under the preferred semantics, there are two extensions: $E_1 = \{\mathtt{a}, \mathtt{c}\}$ modeling the fact that Alice and Carol will attend the party, whereas Bob and David will not; and $E_2 = \{\mathtt{b}, \mathtt{c}\}$ modeling the fact that Bob and Carol will attend the party, whereas Alice and David will not. Observe that the epistemic arguments $\mathtt{c}$ and $\mathtt{d}$ (i.e. the arguments defeated by an epistemic attack) are *deterministic* [7], that is, they have the same acceptance status in all extensions (true for $\mathtt{c}$ and false for $\mathtt{d}$). □

Contributions. We introduce the syntax and the semantics of Epistemic Abstract Argumentation Framework (EAAF) and investigate the complexity of several problems (see below). The proposed EAAF semantics aims to let epistemic arguments be deterministic [7], that is, they have the same acceptance status in all extensions; the status of an argument depends on the credulous or skeptical acceptance of its attackers. Considering the dependence of the status of an argument on its attackers only is inspired by the well-known directionality property proposed for AAF [18, 19], which, if satisfied, then guarantees that the status of each argument depends only on that of its attackers. As we will show in the paper, this allows for the existence of an algorithm for computing EAAF semantics.

Specifically, our main contributions are as follows.

- We formally present EAAF relying on a simple yet expressive form of epistemic attacks, leading to an intuitive epistemic semantics.
- We investigate the complexity of *verification*, *acceptance* and *existence* problems under three well-known semantics: grounded, complete and preferred, all satisfying the above-mentioned directionality property. Our complexity results are summarized in Table 2 (in Section 5). Interestingly, it turns out that the complexity remains the same as that for AAF *i*) for the grounded semantics (irrespective of the considered problem), and *ii*) for the existence problem (irrespective of the considered semantics), while it generally increases w.r.t. that of AAF for the other combinations of considered problem and semantics (verification and acceptance / complete and preferred).
- We propose Algorithm 1 enabling the computation of EAAF semantics at the AAF level. Indeed, the algorithm makes use of an external function that incrementally computes the extensions of an AAF Λ, using the extensions of a sub-AAF Λ' included in Λ. However, the algorithm also works with any external function that computes the extensions of an AAF from scratch, that is, without using previously computed extensions, as done by readily available state-of-the-art AAF solvers [24, 47].

2 PRELIMINARIES

We first review the Dung's framework and then discuss and an extension of AAF with epistemic constraints. Finally, we briefly recall the complexity classes used in the paper.

2.1 Abstract Argumentation Framework

An *Abstract Argumentation Framework* (AAF) is a pair $\langle \mathrm{A}, \Omega \rangle$, where A is a (finite) set of *arguments* and $\Omega \subseteq \mathrm{A} \times \mathrm{A}$ is a set of *attacks* (also called *defeats*). Different argumentation semantics have been proposed for AAF, leading to the characterization of collectively acceptable sets of arguments called *extensions* [33].

Given an AAF $\Lambda = \langle \mathrm{A}, \Omega \rangle$ and a set $S \subseteq \mathrm{A}$ of arguments, an argument $a \in \mathrm{A}$ is said to be *i*) *defeated* w.r.t. S iff $\exists b \in S$ such that $(b, a) \in \Omega$; *ii*) *acceptable* w.r.t. S iff $\forall b \in \mathrm{A}$ with $(b, a) \in \Omega$, $\exists c \in S$ such that $(c, b) \in \Omega$. The sets of defeated and acceptable arguments w.r.t. S are defined as follows (where Λ is understood):

- $\mathit{Def}(S) = \{a \in \mathrm{A} \mid \exists b \in S\,.\,(b, a) \in \Omega\}$;
- $\mathit{Acc}(S) = \{a \in \mathrm{A} \mid \forall b \in \mathrm{A}\,.\,(b, a) \in \Omega \text{ implies } b \in \mathit{Def}(S)\}$.

To simplify the notation, we will often use S^+ to denote $\mathit{Def}(S)$.

Given an AAF $\langle \mathrm{A}, \Omega \rangle$, a set $S \subseteq \mathrm{A}$ of arguments is said to be:

- *conflict-free* iff $S \cap S^+ = \emptyset$;
- *admissible* iff it is conflict-free and $S \subseteq \mathit{Acc}(S)$.

Given an AAF $\langle \mathrm{A}, \Omega \rangle$, a set $S \subseteq \mathrm{A}$ is an *extension* called:

- *complete* ($\mathtt{co}$) iff it is conflict-free and $S = \mathit{Acc}(S)$;
- *preferred* ($\mathtt{pr}$) iff it is a $\subseteq$-maximal complete extension;
- *stable* ($\mathtt{st}$) iff it is a total complete extension, i.e. a complete extension such that $S \cup S^+ = \mathrm{A}$;
- *grounded* ($\mathtt{gr}$) iff it is the $\subseteq$-smallest complete extension.

The set of complete (resp. preferred, stable, grounded) extensions of an AAF Λ will be denoted by $\mathtt{co}(\Lambda)$ (resp. $\mathtt{pr}(\Lambda)$, $\mathtt{st}(\Lambda)$, $\mathtt{gr}(\Lambda)$). It is well-known that the set of complete extensions forms a complete semilattice w.r.t. $\subseteq$, where $\mathtt{gr}(\Lambda)$ is the meet element, whereas the greatest elements are the preferred extensions. All the above-mentioned semantics except the stable admit at least one extension. The grounded semantics, that admits exactly one extension, is said to be a *unique-status* semantics, while the others are said to be *multiple-status* semantics. With a little abuse of notation, in the following we also use $\mathtt{gr}(\Lambda)$ to denote the grounded extension. For any AAF Λ, $\mathtt{st}(\Lambda) \subseteq \mathtt{pr}(\Lambda) \subseteq \mathtt{co}(\Lambda)$ and $\mathtt{gr}(\Lambda) \in \mathtt{co}(\Lambda)$.

Example 3. Let $\Lambda = \langle \mathrm{A}, \Omega \rangle$ be an AAF where $\mathrm{A} = \{\mathtt{a}, \mathtt{b}, \mathtt{c}\}$ and $\Omega = \{(\mathtt{a}, \mathtt{b}), (\mathtt{b}, \mathtt{a}), (\mathtt{b}, \mathtt{c})\}$, whose graph is show in Figure 2 (left). The set of complete extensions of Λ is $\mathtt{co}(\Lambda) = \{E_0 = \emptyset, E_1 = \{\mathtt{a}, \mathtt{c}\}, \mathrm{E}_2 = \{\mathtt{b}\}\}$. E_0 is the grounded extension, while E_1 and E_2 are preferred and stable extensions. □

Given an AAF $\Lambda = \langle \mathrm{A}, \Omega \rangle$ and a semantics $\sigma \in \{\mathtt{gr}, \mathtt{co}, \mathtt{pr}, \mathtt{st}\}$, the *verification* problem (denoted as Ver_σ) is deciding whether a set $S \subseteq \mathrm{A}$ is a σ-extension of Λ. Moreover, for $g \in \mathrm{A}$, the *credulous* (resp. *skeptical*) *acceptance* problem, denoted as CA_σ (resp. SA_σ) is deciding whether g is credulously (resp. skeptically) accepted, that

Figure 2: AAF Λ of Example 3 (left) and AAF Λ' of Example 4 (right).

is deciding whether g belongs to *any* (resp. *every*) σ-extension of Λ. Clearly, $CA_{\texttt{gr}}$ and $SA_{\texttt{gr}}$ coincide.

Recently, a satisfaction problem for AAF called *determinism* (DS_σ) has been introduced [7]. Given a σ-extension E, an argument $g \in \mathrm{A}$ is said to be: *accepted* if $g \in E$; *rejected* if $g \in E$; *undecided* otherwise ($g \notin E \cup E^+$). For a semantic σ, an argument is said to be deterministic if all σ-extensions assign the same status (either accepted, rejected, or undecided) to it.

Finally, the *existence* (resp. *non-empty existence*) problem denoted as Ex_σ (resp. $Ex_\sigma^{\neg\emptyset}$) is deciding whether there exists at least one (resp. at least one non-empty) σ-extension for AAF Λ.

For AAFs, the complexity of the verification, existence and acceptance problems has been investigated (see [35] for an overview). The complexity of the determinism problem is investigated in [7]. The complexity results concerning these problems are summarized in the left-hand side part of Table 2.

Example 4. Consider the AAF Λ of Example 3. Under preferred and stable semantics, both arguments a and b are credulously accepted. None of them is skeptically accepted, nor deterministic.

Considering the AAF Λ' obtained from Λ by adding the self-attack (c, c) (see Figure 2 (right)), there are three complete extensions $E'_0 = \emptyset$, $E'_1 = \{\mathsf{a}\}$ and $E'_2 = \{\mathsf{b}\}$. Both E'_1 and E'_2 are preferred extensions, but only E'_2 is stable. □

An interesting property for argumentation semantics is *directionality* [18, 19]. An argumentation semantics is said to be *directional* if the acceptance status of every argument depends only on the status of its attackers. Thus, the status of an argument is not affected by changes in the status of the arguments that it does not depend on. It turns out that grounded, complete, and preferred semantics are directional, while stable semantics is not [20].

Example 5. Consider the AAFs Λ and Λ' of Example 4. Under stable semantics, argument b that in Λ is not skeptically accepted becomes skeptically accepted in Λ'. In contrast, under preferred semantics, the acceptance status of both a and b does not change. □

2.2 AAF with Epistemic Constraints

An Epistemic Argumentation Framework (EAF) has been proposed in [54]. An EAF is a triple $\langle A, \Omega, C\rangle$, where $\langle A, \Omega\rangle$ is an AAF and C is an epistemic constraint, that is, a propositional formula extended with the modal operators $\mathbf{K}$ and $\mathbf{M}$. Here, the constraint is the belief of an agent which must be satisfied. Intuitively, $\mathbf{K}\phi$ (resp. $\mathbf{M}\phi$) states that the considered agent believes that ϕ is always (resp. possibly) true. EAF semantics is given by sets of feasible extensions of the underlying AAF, called ω-extension sets (ω-labeling sets in [54]), consisting of maximal sets of arguments that satisfies the constraint. There could be different ω-extension sets (ω-sets) for the same epistemic formula, as shown in the following example.

Example 6. Consider the AAF $\Lambda = \langle A = \{\mathsf{a}, \mathsf{b}, \mathsf{c}, \mathsf{d}\}, \Omega = \{(\mathsf{a}, \mathsf{b}), (\mathsf{b}, \mathsf{a}), (\mathsf{c}, \mathsf{d}), (\mathsf{d}, \mathsf{c}), (\mathsf{b}, \mathsf{c})\rangle$ having 5 complete extensions $E_0 = \emptyset$, $E_1 = \{\mathsf{a}\}$, $E_2 = \{\mathsf{a}, \mathsf{c}\}$, $E_3 = \{\mathsf{a}, \mathsf{d}\}$ and $E_4 = \{\mathsf{b}, \mathsf{d}\}$. E_0 is the grounded extension, while E_2, E_3 and E_4 are preferred and stable extensions. Under the preferred semantics, considering the epistemic constraint $C_1 = \mathbf{K}\mathsf{c}$, there exists a unique ω-set $\{E_2\}$ for EAF $\langle A, \Omega, C_1\rangle$, whereas considering $C_2 = \mathbf{K}\mathsf{c} \vee \mathbf{K}\mathsf{d}$ there are the two alternative ω-sets $\{E_2\}$ and $\{E_3, E_4\}$ for EAF $\langle A, \Omega, C_2\rangle$. □

We point out that despite the name Epistemic Argumentation Framework is used, the role of epistemic formulae is only that of introducing constraints over the set of feasible extensions, that is it is similar to that of constraints or preferences in AAF [9, 16, 23, 31].

2.3 Complexity Classes

We recall the main complexity classes used in the paper and, in particular, the definition of the classes P, Σ_h^p, Π_h^p and Δ_h^p, with $h \geq 0$ (see e.g. [50]). P consists of the problems that can be solved in polynomial-time Moreover, we have that:

- $\Sigma_0^p = \Pi_0^p = \Delta_0^p = P$;
- $\Sigma_1^p = NP$ and $\Pi_1^p = coNP$;
- $\Delta_h^p = P^{\Sigma_{h-1}^p}$, $\Sigma_h^p = NP^{\Sigma_{h-1}^p}$, and $\Pi_h^p = co\Sigma_h^p$, $\forall h > 0$.

Thus, P^C (resp. NP^C) denotes the class of problems that can be solved in polynomial time using an oracle in the class C by a deterministic (resp. non-deterministic) Turing machine. The class $\Theta_h^p = \Delta_h^p[log\ n]$ denotes the subclass of Δ_h^p consisting of the problems that can be solved in polynomial time by a deterministic Turing machine performing $O(log\ n)$ calls to an oracle in the class Σ_{h-1}^p.

It is known that:

- $\Sigma_h^p \subset \Delta_{h+1}^p[log\ n] \subset \Delta_{h+1}^p \subset \Sigma_{h+1}^p \subseteq PSPACE$ and
- $\Pi_h^p \subset \Delta_{h+1}^p[log\ n] \subset \Delta_{h+1}^p \subset \Pi_{h+1}^p \subseteq PSPACE$.

3 EPISTEMIC ABSTRACT ARGUMENTATION FRAMEWORK

We augment AAF with epistemic attacks, leading to the concept of Epistemic Abstract Argumentation Framework (EAAF).

3.1 Syntax

We start by introducing the syntax of EAAF.

Definition 1 (Epistemic AAF). *An Epistemic AAF is a quadruple $\Delta = \langle A, \Omega, \Psi, \Phi\rangle$ where A is a set of arguments, $\Omega \subseteq A \times A$ is a set of (standard) attacks, $\Psi \subseteq A \times A$ is a set of weak (epistemic) attacks, and $\Phi \subseteq A \times A$ is a set of strong (epistemic) attacks such that $\Omega \cap \Psi = \Omega \cap \Phi = \Psi \cap \Phi = \emptyset$.*

In the following, we represent attacks $(a, b) \in \Omega$ by $a \rightarrow b$, $(a, b) \in \Psi$ by $a \Rightarrow b$, $(a, b) \in \Phi$ by $a \Rrightarrow b$. An EAAF $\langle A, \Omega, \Psi, \Phi\rangle$ can be seen as a directed graph, where A denotes the set of nodes and Ω, Ψ, and Φ denotes three different kinds of edges. Arguments defeated through epistemic attacks are called *epistemic arguments*.

We say that there is a path from an argument $a \in A$ to argument $b \in A$ if either (i) there exists an attack (a, b) in Δ or (ii) there exists an argument $c \in A$ and two paths, from a to c and from c to b. We say that an argument $b \in A$ *depends* on an argument $a \in A$ if b is *reachable* from a in Δ, that is, if there exists a path from a to b in Δ. Moreover, an argument a depends on attack $\gamma \in (\Omega \cup \Psi \cup \Phi)$ if there exists a path in Δ that contains γ and reaches a.

We now introduce well-formed and plain EAAFs.

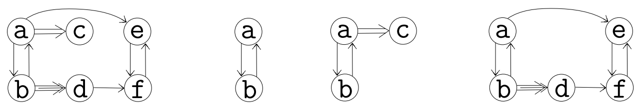

Figure 3: (From left to right) EAAFs Δ, Δ', Δ'' and Δ''' of Example 10.

Definition 2. *An EAAF Δ is said to be:*

- *well-formed if there are no cycles in Δ with epistemic edges.*
- *in plain form if every epistemic argument is attacked by a single (epistemic) attack.*

In the following we assume that our EAAFs are well-formed. The reason for such a restriction is to guarantee that there exists a unique *world view* (c.f. Theorem 1). In the following *we also assume that our EAAFs are in plain form.* There is no loss of generality in making such an assumption as we will show that every well-formed EAAF can be rewritten into an equivalent one in plain form (see Proposition 4). As it will be clear after introducing EAAF semantics, for well-formed EAAFs in plain form, epistemic arguments are deterministic (c.f. Proposition 2).

Example 7. The EAAF of Example 2, $\Delta = \langle A = \{\mathsf{a}, \mathsf{b}, \mathsf{c}, \mathsf{d}\}, \Omega = \{(\mathsf{a}, \mathsf{b}), (\mathsf{b}, \mathsf{a})\}, \Psi = \{(\mathsf{a}, \mathsf{c})\}, \Phi = \{(\mathsf{b}, \mathsf{d})\}\rangle$, whose graph is shown in Figure 1 (right), is well-formed and in plain form. □

The semantics of EAAF is given by relying on the concept of sub-framework (sub-EAAF), which is defined as follows.

Definition 3. *Given two EAAFs Δ and Δ', we say that Δ' is a sub-EAAF of Δ (denoted as $\Delta' \sqsubseteq \Delta$) if Δ' is obtained from Δ by deleting a subset S of the set of epistemic arguments of Δ and all the arguments depending on an argument in S w.r.t. Δ. Moreover, we write $\Delta' \sqsubset \Delta$ if $\Delta' \sqsubseteq \Delta$ and $\Delta' \neq \Delta$.*

Clearly, in Definition 3 by deleting arguments we also delete attacks having as a source or target element a deleted argument.

Example 8. Consider the EAAF $\Delta = \langle\{\mathsf{a}, \mathsf{b}, \mathsf{c}, \mathsf{d}, \mathsf{e}, \mathsf{f}\}, \{(\mathsf{a}, \mathsf{b}), (\mathsf{b}, \mathsf{a}), (\mathsf{a}, \mathsf{e}), (\mathsf{d}, \mathsf{f}), (\mathsf{e}, \mathsf{f}), (\mathsf{f}, \mathsf{e})\}, \{(\mathsf{a}, \mathsf{c})\}, \{(\mathsf{b}, \mathsf{d})\}\rangle$ shown in Figure 3 (left). We have four sub-EAAFs $\Delta^* \sqsubseteq \Delta$, as shown in the figure: the first one (from left to right) coincides with Δ, the others are obtained by deleting all arguments depending on: (i) both arguments c and d, (ii) only d, and (iii) only c, respectively. □

3.2 Semantics

We first introduce the semantics of EAAF and then present some results concerning properties of the proposed framework.

For any EAAF $\Delta = \langle A, \Omega, \Psi, \Phi\rangle$, a set W of sets of arguments in A is called *world view* of Δ. Informally, a world view can be seen as a set of extensions that are to be used to compute the status of epistemic arguments. Given EAAF $\Delta' = \langle A', \Omega', \Psi', \Phi'\rangle \sqsubseteq \Delta$, we denote by $W_{\downarrow\Delta'} = \{S \cap A' \mid S \in W\}$ the projection of W over A'.

With the aim of providing EAAF semantics by extending AF semantics, we first extend the definitions of defeated and acceptable arguments for EAAF by taking into account the additional concept of world view, that is a candidate set of extensions, which is used to decide if an argument is epistemically defeated/acceptable. Given an EAAF Δ, a world view W of Δ, and a set $S \in W$, the sets of arguments defeated (resp. accepted) w.r.t. S and W are defined as follows:

- $Def(W, S) = \{b \in A \mid (\exists a \in S \,.\, a \rightarrow b) \lor$
  $(\exists T \in W \,.\, \exists a \in T \,.\, a \Rrightarrow b) \lor$
  $(\forall T \in W \,.\, \exists a \in T \,.\, a \Rightarrow b)\}$.
- $Acc(W, S) = \{b \in A \mid \forall a \in A\,.$
  $((a \rightarrow b)$ implies $a \in Def(W, S)) \land$
  $((a \Rrightarrow b)$ implies $\forall T \in W \,.\, a \in Def(W, T)) \land$
  $((a \Rightarrow b)$ implies $\exists T \in W \,.\, a \in Def(W, T))$.

Example 9. Considering the EAAF Δ of Example 8 and the world view $W = \{S_1 = \{\mathsf{c}\}, S_2 = \{\mathsf{a}, \mathsf{c}\}, S_3 = \{\mathsf{b}, \mathsf{c}\}\}$, we have that:

- $Def(W, S_1) = \{\mathsf{d}\}$ and $Acc(W, S_1) = \{\mathsf{c}\}$;
- $Def(W, S_2) = \{\mathsf{b}, \mathsf{d}\}$ and $Acc(W, S_2) = \{\mathsf{a}, \mathsf{c}\}$; and
- $Def(W, S_3) = \{\mathsf{a}, \mathsf{d}\}$ and $Acc(W, S_3) = \{\mathsf{b}, \mathsf{c}\}$. □

Given an EAAF Δ and a world view W of Δ, a set $S \in W$ is:

- *W-conflict-free* if $S \cap Def(W, S) = \emptyset$;
- *W-admissible* if it is W-conflict-free and $S \subseteq Acc(W, S)$;
- *W-complete* (W-co) if it is W-conflict-free and $S = Acc(W, S)$.

Moreover, a W-complete set S is said to be *W-preferred* (W-pr) if S is $\subseteq$-maximal, and *W-grounded* (W-gr) if S is $\subseteq$-minimal.

We are now ready to define EAAF semantics.

Definition 4 (EAAF Semantics). *Let $\sigma \in \{\mathsf{gr}, \mathsf{co}, \mathsf{pr}\}$ be a semantics and W a world view of EAAF Δ. Then, W is a σ-world view for Δ if $\forall \Delta' \sqsubseteq \Delta$ the following conditions hold:*

(i) *every $S \in W_{\downarrow\Delta'}$ is a $W_{\downarrow\Delta'}$-σ set, and*
(ii) *there is no world view W^* for Δ' such that $W_{\downarrow\Delta'} \subset W^*$ and every $S^* \in W^*$ is W^*-σ for Δ'.*

We now explain Definition 4. Given a semantics σ, a W-σ set intuitively represents a candidate set of σ-extensions for an EAAF. Then, such a set turns out to actually be a set of extensions if the conditions in Definition 4 hold, whose rationale is as follows. Given a world view W of an EAAF Δ, we check that for all sub-frameworks Δ', every element $S \in W' = W_{\downarrow\Delta'}$ is a W'-σ set (condition i) and W' is maximal (condition ii). Intuitively, the first condition ensures that the status of an argument is confirmed in all sub-frameworks considered. This is due to the fact that, focusing on semantics that are directional, the status of every argument a has to be 'confirmed' also in all sub-frameworks Δ' of Δ containing a. This condition implicitly holds for all AF semantics that are directional (e.g. grounded, complete, preferred). The second condition of Definition 4 ensures that, if there is a larger σ-world view for which condition i) holds, then we prefer to take it. That is, intuitively, we aim at having the whole set of extensions. We show that this set is unique in Theorem 1. Finally, as shown below in Example 10, checking that the above-mentioned conditions hold for all sub-frameworks is important to avoid returning wrong conclusions (i.e., world views that contradict our intuition).

It is worth noting that whenever $\Psi = \Phi = \emptyset$, we have that the definitions of defeated and acceptable arguments coincide with the ones defined for AAF, that is $Def(\{S\}, S) = Def(S)$ and $Acc(\{S\}, S) = Acc(S)$. This lead to the following result that states that EAAF semantics extends that of AAF.

Proposition 1. *Let $\Delta = \langle A, \Omega, \Psi, \Phi\rangle$ be a well-formed EAAF with $\Psi = \Phi = \emptyset$, $\sigma \in \{\mathsf{gr}, \mathsf{co}, \mathsf{pr}\}$, and $\Lambda = \langle A, \Omega\rangle$ the AAF corresponding to Δ. Then, Δ has a unique σ-world view $W = \sigma(\Lambda)$.*

Table 1: σ-world view for each EAAF $\Delta^* \sqsubseteq \Delta$ in Figure 3.

Δ^*	$\mathtt{gr}(\Delta^*)$	$\mathtt{co}(\Delta^*)$	$\mathtt{pr}(\Delta^*)$
Δ'	$\{\emptyset\}$	$\{\emptyset, \{a\}, \{b\}\}$	$\{\{a\}, \{b\}\}$
Δ''	$\{\emptyset\}$	$\{\{c\}, \{a, c\}, \{b, c\}\}$	$\{\{a, c\}, \{b, c\}\}$
Δ'''	$\{\emptyset\}$	$\{\emptyset, \{f\}, \{a, f\}, \{b\}, \{b, e\}, \{b, f\}\}$	$\{\{a, f\}, \{b, e\}, \{b, f\}\}$
Δ	$\{\emptyset\}$	$\{\{c\}, \{c, f\}, \{a, c, f\}, \{b, c\}, \{b, c, e\}, \{b, c, f\}\}$	$\{\{a, c, f\}, \{b, c, e\}, \{b, c, f\}\}$

As stated next, any well-formed EAAF has a unique world view.

Theorem 1. *Any well-formed EAAF has a unique σ-world view, with $\sigma \in \{\mathtt{gr}, \mathtt{co}, \mathtt{pr}\}$.*

For any (well-formed) EAAF Δ and semantics $\sigma \in \{\mathtt{gr}, \mathtt{co}, \mathtt{pr}\}$ we use $\sigma(\Delta)$ to denote the σ-world view of Δ, and will often call its elements σ-extensions.

Example 10. Continuing with Example 8, Table 1 reports the σ-world view for each EAAF $\Delta^* \sqsubseteq \Delta$ in Figure 3 and $\sigma \in \{\mathtt{gr}, \mathtt{co}, \mathtt{pr}\}$.

Now, consider the EAAF Δ'' (shown in Figure 3), the world view $W = \{S = \{a\}\}$, and the preferred semantics. If in Definition 4 we had only focused on the given EAAF Δ'' without looking at its sub-frameworks, as S is a W-preferred set and W is maximal (i.e., both conditions (*i*) and (*ii*) of Definition 4 are satisfied if focusing on Δ'' only), we would have concluded that c is defeated. However, we had expected that c would have been accepted. Indeed, according to Definition 4, the only preferred-world view of Δ'' is $W'' = \{\{a, c\}, \{b, c\}\}$ (cf. Table 1). In fact, considering the sub-framework Δ' (cf. Figure 3) obtained from Δ'' by deleting the epistemic argument c, the only preferred world view of Δ' is $W' = W''_{\downarrow\Delta'} = \{\{a\}, \{b\}\}$, which using Definition 4 allows us to discard $W = \{\{a\}\}$ from being a preferred-world view of Δ''. □

Although Definition 4 requires to check a (possibly) exponential number of sub-EAAF Δ' of a given EAAF Δ, in the next section we show that only a linear number (w.r.t. $|(\Psi \cup \Phi)|$) of EAAFs $\Delta' \sqsubseteq \Delta$ needs to be considered.

According to the proposed EAAF semantics, epistemic arguments are deterministic, that is, they have the same "truth assignment" in a world view, that in turn depends on either the credulous or skeptical acceptance of its attackers.

Proposition 2. *Let $\Delta = \langle A, \Omega, \Psi, \Phi \rangle$ be an EAAF, $\sigma \in \{\mathtt{gr}, \mathtt{co}, \mathtt{pr}\}$ a semantics, and W the σ-world view of Δ. Then, any epistemic argument $x \in A$ is deterministic, that is, one of the following three conditions hold:*

i) $\forall\, S \in W .\ x \in Acc(W, S)$; *or*
ii) $\forall\, S \in W .\ x \in Def(W, S)$; *or*
iii) $\forall\, S \in W .\ x \notin (Acc(W, S) \cup Def(W, S))$.

This property is related to the directionality property for AAF [18, 19], that guarantees that the status of each argument depends only on that of its attackers. As we will see in Section 4, this leads to the existence of an algorithm for computing EAAF semantics (i.e., Algorithm 1). For this reason and the fact that stable semantics does not guarantee the existence of at least one extension, we will consider the definition of stable semantics for EAAF in future work.

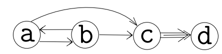

Figure 4: EAAF Δ of Example 11.

An alternative way to define preferred (resp. grounded) extensions for EAAF could be that of choosing among complete extensions those that are maximal (resp. minimal) w.r.t set-inclusion, as it is done for AAF. More in detail, given an EAAF Δ and its complete-world view $W = \mathtt{co}(\Delta)$, we could have defined the preferred-world (resp. grounded-world) view for Δ as $\mathtt{pr}(\Delta) = \{S \in W \mid S \text{ is maximal w.r.t. } \subseteq\}$ (resp. $\mathtt{gr}(\Delta) = \{S \in W \mid S \text{ is minimal w.r.t. } \subseteq\}$). This is different from what is done in Definition 4 where to define a pr-world view (resp. gr-world view) we start with a world view W that is not necessarily $\mathtt{co}(\Delta)$. However, the above-mentioned alternative way to define preferred extensions for EAAF may lead to counter-intuitive solutions, as shown in the following example.

Example 11. Consider the EAAF $\Delta = \langle \{a, b, c, d\}, \{(a, b), (b, a), (a, c), (b, c)\}, \{(c, d)\}, \emptyset \rangle$, shown in Figure 4, and the preferred semantics. Intuitively, the strong epistemic attack states that d is accepted if c is skeptically rejected. The preferred extensions of Δ, that is, the elements in its pr-world view are $\{a, d\}$ and $\{b, d\}$. Thus, we obtain that c is skeptically defeated and, consequently, d is accepted.

However, if we start with the complete-world view $\mathtt{co}(\Delta)$, we have that there are three complete extensions $S_1 = \emptyset$, $S_2 = \{a\}$ (with b and c defeated and d undecided) and $S_3 = \{b\}$ (with a and c defeated and d undecided). As the $\subseteq$-maximal sets in $\mathtt{co}(\Delta)$ are S_2 and S_3, we conclude that under the above-mentioned "alternative" preferred semantics d is undecided, whereas c is (skeptically) false, contradicting our intuition. □

As stated next, differently from AAF, grounded and preferred extensions are not guaranteed to be complete extensions of EAAF. Related to this, even in AAF credulous and skeptical acceptance may give different results under different semantics.

Proposition 3.

- *There exists an EAAF Δ s.t. $S \in \mathtt{pr}(\Delta) \wedge S \notin \mathtt{co}(\Delta)$.*
- *There exists an EAAF Δ s.t. $S \in \mathtt{gr}(\Delta) \wedge S \notin \mathtt{co}(\Delta)$.*

Particularly, for the first item, consider the EAAF $\Delta = \langle \{a, b, c, d, e, f\}, \{(a, b), (b, a), (a, c), (b, c), (c, d)\}, \{(d, e)\ (e, f)\}, \emptyset \rangle$. With a little effort, it ca be checked that $\mathtt{pr}(\Delta) = \{S_1 = \{a, d, f\}, S_2 = \{b, d, f\}\}$ and $\mathtt{co}(\Delta) = \{\emptyset, \{a, d\}, \{b, d\}\}$, and thus neither $S_1 \in \mathtt{co}(\Delta)$ nor $S_2 \in \mathtt{co}(\Delta)$. As for the second item, consider the EAAF Δ of Example 10, it holds that $\mathtt{gr}(\Delta) = \emptyset$ and $\emptyset \notin \mathtt{co}(\Delta)$.

4 COMPUTATION

We first show that there is no loss of generality in making the assumption that EAAFs are in plain form. In fact, any EAAF Δ can be rewritten into an "equivalent" one $\widehat{\Delta}$ in plain form such that complete (resp. grounded, preferred) world view of Δ can be obtained from that of $\widehat{\Delta}$, as stated in the following proposition.

Proposition 4. *Let Δ be an EAAF and W the σ-world view for Δ, with $\sigma \in \{\mathtt{gr}, \mathtt{co}, \mathtt{pr}\}$. Then, there exists an EAAF $\widehat{\Delta}$ in plain form whose σ-world view $\widehat{W}$ is such that $W = \widehat{W}_{\downarrow\Delta}$ holds.*

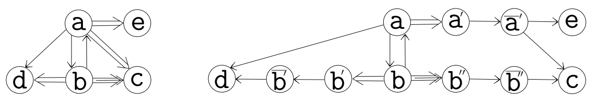

Figure 5: (From left to right) EAFs Δ and its plain version $\widehat{\Delta}$ of Example 12.

An example of the application of Proposition 4 is the following.

Example 12. Consider the EAAF $\Delta = \langle A = \{\mathsf{a}, \mathsf{b}, \mathsf{c}, \mathsf{d}, \mathsf{e}\}, \Omega = \{(\mathsf{a}, \mathsf{b}), (\mathsf{b}, \mathsf{a}), (\mathsf{a}, \mathsf{d})\}, \Psi = \{(\mathsf{a}, \mathsf{c}), (\mathsf{a}, \mathsf{e}), (\mathsf{b}, \mathsf{d})\}, \Phi = \{(\mathsf{b}, \mathsf{c})\}$, shown in Figure 5 (left), and its preferred-world view $W = \{\{\mathsf{a}, \mathsf{e}\}, \{\mathsf{b}, \mathsf{d}, \mathsf{e}\}\}$. Let $\widehat{\Delta} = \langle \widehat{A}, \widehat{\Omega}, \widehat{\Psi}, \widehat{\Phi} \rangle$, shown in Figure 5 (right), where:

- $\widehat{A} = \{\mathsf{a}, \mathsf{b}, \mathsf{c}, \mathsf{d}, \mathsf{e}, \mathsf{a}', \overline{\mathsf{a}}', \mathsf{b}', \overline{\mathsf{b}}', \mathsf{b}'', \overline{\mathsf{b}}''\}$;
- $\widehat{\Omega} = \{(\mathsf{a}, \mathsf{b}), (\mathsf{b}, \mathsf{a}), (\mathsf{a}, \mathsf{d}), (\mathsf{a}', \overline{\mathsf{a}}'), (\overline{\mathsf{a}}', \mathsf{c}), (\overline{\mathsf{a}}', \mathsf{e}), (\mathsf{b}', \overline{\mathsf{b}}'), (\overline{\mathsf{b}}', \mathsf{d}), (\mathsf{b}'', \overline{\mathsf{b}}''), (\overline{\mathsf{b}}'', \mathsf{c})\}$;
- $\widehat{\Psi} = \{(\mathsf{a}, \overline{\mathsf{a}}'), (\mathsf{b}, \overline{\mathsf{b}}')\}$; and
- $\widehat{\Phi} = \{(\mathsf{b}, \overline{\mathsf{b}}'')\}$.

The preferred-world view for $\widehat{\Delta}$ is $\widehat{W} = \{\{\mathsf{a}, \mathsf{e}, \mathsf{a}', \mathsf{b}', \overline{\mathsf{b}}''\}, \{\mathsf{b}, \mathsf{d}, \mathsf{e}, \mathsf{a}', \mathsf{b}', \overline{\mathsf{b}}''\}\}$. As prescribed by Proposition 4, $W = \widehat{W}_{\downarrow\Delta}$. □

Thus, as previously stated in Section 3.1, w.l.o.g. we can assume that EAFs are in plain form.

For any EAAF $\Delta = \langle A, \Omega, \Psi, \Phi \rangle$, we use $\Gamma = (\Psi \cup \Phi) = \{\gamma_1 = (a_1, b_1), \dots, \gamma_n = (a_n, b_n)\}$ to denote the set of epistemic attacks in Δ. Observe that, given Γ, we can define a partial order PO over it such that for every $\gamma_i = (a_i, b_i) \in \Gamma$ and $\gamma_j = (a_j, b_j) \in \Gamma$ with $i < j$, it holds that a_i does not depends on b_j. Moreover, we use $\overline{\Gamma} = \langle \gamma_1, \dots, \gamma_m \rangle$ to denote a linear ordering over Γ (compatible with PO). Clearly, we may have different linear orderings over Γ, as shown in the following example.

Example 13. Consider the EAAF $\Delta = \langle \{\mathsf{a}, \mathsf{b}, \mathsf{c}, \mathsf{d}, \mathsf{e}, \mathsf{f}\}, \{(\mathsf{a}, \mathsf{b}), (\mathsf{b}, \mathsf{a}), (\mathsf{d}, \mathsf{f}), (\mathsf{e}, \mathsf{f})\}, \{\gamma_1, \gamma_3\}, \{\gamma_2\} \rangle$ where $\gamma_1 = (\mathsf{a}, \mathsf{c})$, $\gamma_2 = (\mathsf{c}, \mathsf{e})$, and $\gamma_3 = (\mathsf{b}, \mathsf{d})$. The partial order PO over $\{\gamma_1, \gamma_2, \gamma_3\}$ states that γ_1 *precedes* γ_2. We have three linear orderings over Γ, compatible with PO, that are $\langle \gamma_3, \gamma_1, \gamma_2 \rangle$, $\langle \gamma_1, \gamma_2, \gamma_3 \rangle$, and $\langle \gamma_1, \gamma_3, \gamma_2 \rangle$. □

An Algorithm for Computing EAAF Semantics. Although Definition 4 requires to check a (possibly) exponential number of EAFs $\Delta' \sqsubseteq \Delta$, we show that only a linear number of EAAFs needs to be considered. This is achieved by providing an algorithm (i.e., Algorithm 1) that is able to compute the σ-world view of an EAAF $\Delta = \langle A, \Omega, \Psi, \Phi \rangle$, with $\sigma \in \{\mathsf{gr}, \mathsf{co}, \mathsf{pr}\}$, by following any of the linear orderings on $\Gamma = \Psi \cup \Phi$. Before presenting Algorithm 1 we introduce some notation.

Let $\Delta = \langle A, \Omega, \Psi, \Phi \rangle$ be EAAF and $\overline{\Gamma} = \langle \gamma_1, \dots, \gamma_n \rangle$ a linear ordering over $\Gamma = \Psi \cup \Phi$, and $\Delta_i \sqsubseteq \Delta$ (with $i \in [0, n]$) the sub-EAAF containing only the arguments not depending on $\gamma_{i+1}, \dots, \gamma_n$. In particular, $\Delta_0 = \langle A, \Omega, \emptyset, \emptyset \rangle$ is the sub-EAAF containing no epistemic arguments, Δ_1 is the sub-EAAF containing only the epistemic attack γ_1, and so on up to $\Delta_n = \Delta$.

Definition 5 (Eval Function). *Given an AAF $\Lambda = \langle A, \Omega \rangle$, a semantics $\sigma \in \{\mathsf{gr}, \mathsf{co}, \mathsf{pr}\}$, the set of σ-extensions of Λ $\sigma(\Lambda)$, and an epistemic attack (a, b) s.t. $a \in A$ and $b \notin A$, the function $\mathsf{eval}_\sigma((a, b), \Lambda)$ computes the truth value of b over Λ w.r.t. σ as follows:*

$$
\mathsf{eval}_\sigma((a,b),\Lambda) = \begin{cases} \mathsf{F} & \text{if } \; a \Rightarrow b \wedge \forall E \in \sigma(\Lambda) \,.\, a \in E \,\vee \\ & \quad\; a \Rrightarrow b \wedge \exists E \in \sigma(\Lambda) \,.\, a \in E \\ \mathsf{T} & \text{if } \; a \Rightarrow b \wedge \exists E \in \sigma(\Lambda) \,.\, a \in E^+ \,\vee \\ & \quad\; a \Rrightarrow b \wedge \forall E \in \sigma(\Lambda) \,.\, a \in E^+ \,\vee \\ \mathsf{U} & \text{otherwise.} \end{cases}
$$

Intuitively, $\mathsf{eval}_\sigma((a, b), \Lambda)$ assigns to the epistemic argument b a truth value w.r.t. the σ-extensions of Λ on the basis of the meaning of epistemic attack (a, b).

Algorithm 1 SolveEAAF

Require:
A (well formed) EAAF $\Delta = \langle A, \Omega, \Psi, \Phi \rangle$;
A linear ordering $\overline{\Gamma} = \langle (a_1, b_1), \dots, (a_n, b_n) \rangle$ over $\Gamma = (\Psi \cup \Phi)$;
A semantics $\sigma \in \{\mathsf{gr}, \mathsf{co}, \mathsf{pr}\}$; and
A (sound and complete) function $\mathsf{IncrSolve}_\sigma(\Lambda, \cdot)$ that incrementally returns the set of σ-extensions of an AAF Λ.
Ensure: The σ-world view W of Δ
1: Compute $\Delta_i = \langle A_i, \Omega_i, \Psi_i, \Phi_i \rangle$ for $i \in [0, n]$;
2: Let $\delta_i = \langle A_i^\delta = A_i \setminus A_{i-1}, \Omega_i^\delta = \Omega \cap (A_i^\delta \times A_i^\delta) \rangle$ for $i \in [1, n]$;
3: Let $\Lambda_0 = \langle A_0^\Lambda = A_0, \Omega_0^\Lambda = \Omega_0 \rangle$;
4: $W_0 = \mathsf{IncrSolve}_\sigma(\Lambda_0, \emptyset)$;
5: **for** $i \in [1, n]$ **do**
6: Let $\Lambda_i = \langle A_i^\Lambda, \Omega_i^\Lambda \rangle$ with
$A_i^\Lambda = A_{i-1}^\Lambda \cup A_i^\delta \setminus \{b_i \mid \mathsf{eval}_\sigma((a_i, b_i), \Lambda_{i-1}) = \mathsf{F}\}$,
$\Omega_i^\Lambda = \Omega_{i-1}^\Lambda \cup (\Omega \cap (A_i^\Lambda \times A_i^\delta)) \cup$
$\{(b_i, b_i) \mid \mathsf{eval}_\sigma((a_i, b_i), \Lambda_{i-1}) = \mathsf{U}\}$;
7: $W_i = \mathsf{IncrSolve}_\sigma(\Lambda_i, W_{i-1})$;
return $W = W_n$;

We are now ready to present Algorithm 1 that computes the (unique) world view of the input EAAF Δ under semantics σ.

The algorithm makes use of an external function IncrSolve that computes the σ-extensions of an AAF Λ, knowing the extensions of a sub-AAF Λ' of Λ.

Algorithm 1 works as follows. The $n+1$ EAAFs Δ_i (with $i \in [0, n]$) containing only the arguments not depending on $\gamma_{i+1}, \dots, \gamma_n$ are computed at line 1. Then, n AAFs $\delta_i = \langle A_i^\delta, \Omega_i^\delta \rangle$ (with $i \in [1, n]$) are computed at line 2. Intuitively, δ_i contains the arguments and attacks that are in Δ_i but not in Δ_{i-1}. In other words, δ_i is obtained from Δ_i after removing γ_i and all arguments of Δ_{i-1}, and represents the part of Δ_i that depends on γ_i. Then, at line 3 the AAF Λ_0 is computed. It consists of elements of Δ that do not depend on any epistemic attack, and thus coincides with Δ_0. The set of σ-extensions W_0 of Λ_0 is then computed at line 4 by calling the function $\mathsf{IncrSolve}_\sigma(\Lambda_0, \emptyset)$, where $\emptyset$ means that no information (i.e. extensions previously computed) is exploited during the computation. Then, for each step $i \in [1, n]$, Algorithm 1 (incrementally) computes:

(*a*) the AAF Λ_i that extends the previous one (i.e. Λ_{i-1}) by adding:
- all the non-epistemic arguments of δ_i, plus the epistemic argument b_i if it is not evaluated to be false by function eval_σ;
- the (standard) attacks occurring in Δ_i towards arguments in δ_i;
- a self (standard) attack (b_i, b_i) if b_i is evaluated to be undecided by function eval_σ;

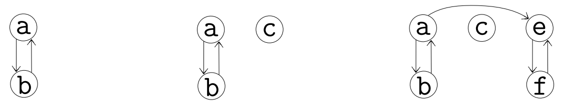

Figure 6: (From left to right) AAFs Λ_0, Λ_1 and Λ_2 of Example 14.

(b) the set of σ-extensions W_i of Λ_i by calling the function $\mathsf{IncrSolve}_\sigma(\Lambda_i, W_{i-1})$ that, given the world view $W_{i-1} = \sigma(\Delta_{i-1})$ for $\Delta_{i-1} \sqsubset \Delta_i$, computes the σ-extensions of the AAF Λ_i (derived from Δ_i), by exploiting the extensions in W_{i-1}, thus avoiding wasted effort in the computation.[1]

Function eval_σ evaluates at each step i the epistemic argument b_i on the basis of AAF Λ_{i-1} and the set of its σ-extensions previously computed, which intuitively contain the information needed to evaluate b_i. The algorithm ends by returning $W = W_n$ at line 7.

An example of the execution of Algorithm 1 is as follows.

Example 14. Consider the EAAF $\Delta = \langle A, \Omega, \Psi, \Phi\rangle$ of Example 8 (see Figure 3 (left)), the semantics $\sigma = \mathsf{pr}$, and the linear ordering $\overline{\Gamma} = (\gamma_1 = (\mathsf{a}, \mathsf{c}), \gamma_2 = (\mathsf{b}, \mathsf{d}))$ of $\Gamma = \Psi \cup \Phi$.

The three EAAFs Δ_i for $i \in [0, 2]$, computed at line 1, are:

- $\Delta_0 = \langle A_0 = \{\mathsf{a}, \mathsf{b}\}, \Omega_0 = \{(\mathsf{a}, \mathsf{b}), (\mathsf{b}, \mathsf{a})\}, \Psi_0 = \emptyset, \Phi_0 = \emptyset\rangle$, corresponding to EAAF Δ' in Figure 3;
- $\Delta_1 = \langle A_1 = \{\mathsf{a}, \mathsf{b}, \mathsf{c}\}, \Omega_1 = \{(\mathsf{a}, \mathsf{b}), (\mathsf{b}, \mathsf{a})\}, \Psi_1 = \{(\mathsf{a}, \mathsf{c})\}, \Phi_1 = \emptyset\rangle$, corresponding to EAAF Δ'' in Figure 3;
- $\Delta_2 = \langle A_2 = A, \Omega_2 = \Omega, \Psi_2 = \Psi, \Phi_2 = \Phi\rangle$, corresponding to EAAF Δ in Figure 3 (left).

The two AAFs computed at line 2 are:

- $\delta_1 = \langle A_1^\delta = \{\mathsf{c}\}, \Omega_1^\delta = \emptyset\rangle$, and
- $\delta_2 = \langle A_2^\delta = \{\mathsf{d}, \mathsf{e}, \mathsf{f}\}, \Omega_2^\delta = \{(\mathsf{d}, \mathsf{f}), (\mathsf{e}, \mathsf{f}), (\mathsf{f}, \mathsf{e})\}\rangle$.

Then, $\Lambda_0 = \langle A_0^\Lambda = A_0, \Omega_0^\Lambda = \Omega_0\rangle$ is computed at line 3, and the set of preferred extension $W_0 = \{\{\mathsf{a}\}, \{\mathsf{b}\}\}$ is computed at line 4 by calling the external AAF solver $\mathsf{IncrSolve}_{\mathsf{pr}}(\Lambda_0, \emptyset)$.

Then, at line 6 the algorithm computes:

- for $i = 1$, the AAF $\Lambda_1 = \langle A_1^\Lambda, \Omega_1^\Lambda\rangle$ (shown in Figure 6, center) with $A_1^\Lambda = \{\mathsf{a}, \mathsf{b}\} \cup (\{\mathsf{c}\} \setminus \emptyset)$ and $\Omega_1^\Lambda = \{(\mathsf{a}, \mathsf{b}), (\mathsf{b}, \mathsf{a})\}$, as $\mathsf{eval}_\sigma((a_1 = \mathsf{a}, b_1 = \mathsf{c}), \Lambda_0) = \mathsf{T}$; also, the set of its preferred extensions is $W_1 = \{\{\mathsf{a}, \mathsf{c}\}, \{\mathsf{b}, \mathsf{c}\}\}$;
- for $i = 2$, the AAF $\Lambda_2 = \langle A_2^\Lambda, \Omega_2^\Lambda\rangle$ (shown in Figure 6, right) with $A_2^\Lambda = \{\mathsf{a}, \mathsf{b}, \mathsf{c}\} \cup (\{\mathsf{d}, \mathsf{e}, \mathsf{f}\} \setminus \{\mathsf{d}\})$ and $\Omega_2^\Lambda = \{(\mathsf{a}, \mathsf{b}), (\mathsf{b}, \mathsf{a})\} \cup \{(\mathsf{a}, \mathsf{e}), (\mathsf{e}, \mathsf{f}), (\mathsf{f}, \mathsf{e})\} \cup \emptyset$, as $\mathsf{eval}_\sigma((a_2 = \mathsf{b}, b_2 = \mathsf{d}), \Lambda_1) = \mathsf{F}$; the set of its preferred extensions is $W_2 = \{\{\mathsf{a}, \mathsf{c}, \mathsf{f}\}, \{\mathsf{b}, \mathsf{c}, \mathsf{e}\}, \{\mathsf{b}, \mathsf{c}, \mathsf{f}\}\}$.

Algorithm 1 ends by returning the set of preferred extensions of Δ, that is $\mathsf{pr}(\Delta) = W_2 = \{\{\mathsf{a}, \mathsf{c}, \mathsf{f}\}, \{\mathsf{b}, \mathsf{c}, \mathsf{e}\}, \{\mathsf{b}, \mathsf{c}, \mathsf{f}\}\}$. □

As stated next, Algorithm 1 is sound and complete, independently of the chosen linear ordering over the set of epistemic attacks.

Theorem 2. *Let $\Delta = \langle A, \Omega, \Psi, \Phi\rangle$ be an EAAF, $\overline{\Gamma}$ a linear ordering over $\Gamma = (\Psi \cup \Phi)$, and $\sigma \in \{\mathsf{gr}, \mathsf{co}, \mathsf{pr}\}$ a semantics. Then, Algorithm 1 returns the σ-world view of Δ.*

[1] In SAT-based solvers [24, 47], this can be achieved by introducing additional clauses encoding the status of arguments of the previous extensions. Approaches for the incremental computation in argumentation have been investigated in e.g. [2–4].

5 COMPLEXITY

We investigate the complexity of fundamental reasoning problems for EAAF. In particular, we study the *verification*, *existence*, and *credulous/skeptical* acceptance problems, that are usually considered for analyzing the complexity of argumentation frameworks.

Given an EAAF $\Delta = \langle A, \Omega, \Psi, \Phi\rangle$ and a semantics $\sigma \in \{\mathsf{gr}, \mathsf{co}, \mathsf{pr}\}$:

- the *verification* problem for EAAF, denoted as VE_σ, consists in deciding whether a given set of arguments $S \subseteq A$ is a σ-extension of Δ, that is, deciding whether S is in the σ-world view of Δ;
- the *existence* (resp. *non-empty existence*) problem for EAAF, denoted as Ex_σ (resp. $Ex_\sigma^{\neg\emptyset}$) with $\sigma \in \{\mathsf{gr}, \mathsf{co}, \mathsf{pr}\}$ consists in deciding whether there exists at least one (resp. at least one non-empty) σ-extension S for Δ;
- the *credulous* (resp. *skeptical*) acceptance problem, denoted as CA_σ (resp. SA_σ), consists in deciding whether a given goal argument $g \in A$ belongs to any (resp. every) σ-extension of Δ.

Observe that if argument g is epistemic, credulous and skeptical acceptance problems coincide (cf. Proposition 2). Therefore, we call simply this problem *epistemic acceptance* and denote it as EA_σ.

The following fact states that the epistemic acceptance problem captures the credulous and skeptical acceptance problems also for non-epistemic arguments.

Fact 1. *Let $\Delta = \langle A, \Omega, \Psi, \Phi\rangle$ be an EAAF, $g \in A$ any of its non-epistemic arguments, and $\sigma \in \{\mathsf{gr}, \mathsf{co}, \mathsf{pr}\}$ a semantics. Then:*

- *$CA_\sigma(\Delta, g) = EA_\sigma(\Delta', g'')$ with :*
  $\Delta' = \langle A \cup \{g', g''\}, \Omega \cup \{(g, g')\}, \Psi \cup \{(g', g'')\}, \Phi\rangle$
- *$SA_\sigma(\Delta, g) = EA_\sigma(\Delta', g'')$ with :*
  $\Delta' = \langle A \cup \{g', g''\}, \Omega \cup \{(g, g')\}, \Psi, \Phi \cup \{(g', g'')\}\rangle$

Thus, asking for the credulous and skeptical acceptance of an argument g w.r.t. an EAAF Δ is equivalent to asking for the epistemic acceptance of a fresh epistemic argument g'' w.r.t. an EAAF Δ', that is obtained from Δ by adding only a pair of attacks.

For this reason and for the fact that epistemic arguments are deterministic (Proposition 2), w.l.o.g. we study the complexity of verification, existence, epistemic acceptance problems in EAAFs (without considering credulous and skeptical acceptance that, as shown above, can be immediately reduced to epistemic acceptance).

Theorem 3. *VE_σ is: i) in P for $\sigma = \mathsf{gr}$, ii) Θ_2^p-hard and in Δ_2^p for $\sigma = \mathsf{co}$, and iii) Θ_2^p-hard and in Δ_3^p for $\sigma = \mathsf{pr}$.*

The next theorem states the complexity of epistemic acceptance.

Theorem 4. *EA_σ is: i) in P for $\sigma = \mathsf{gr}$, ii) Θ_2^p-hard and in Δ_2^p for $\sigma = \mathsf{co}$, and iii) Π_2^P-h and in Δ_3^P for $\sigma = \mathsf{pr}$.*

The following corollary states that for EAAF there always exists at least one extension, as for the case of AAF.

Corollary 1. *Ex_σ is trivial for any $\sigma \in \{\mathsf{gr}, \mathsf{co}, \mathsf{pr}\}$.*

Finally, deciding the existence of a non-empty extension has the same complexity as that for AAF.

Theorem 5. *$Ex_\sigma^{\neg\emptyset}$ is: i) in P for $\sigma = \mathsf{gr}$, and ii) NP-complete for $\sigma \in \{\mathsf{co}, \mathsf{pr}\}$.*

The results of this section, along with some related complexity results for AAF, are summarized in Table 2. Although the complexity

Table 2: Complexity of the verification (Ver_σ), credulous acceptance (CA_σ), skeptical acceptance (SA_σ), existence (Ex_σ), non-empty existence ($Ex_\sigma^{\neg\emptyset}$), and determinism problems for AAF, and of the verification (VE_σ), epistemic acceptance (EA_σ), existence (Ex_σ), and non-empty existence ($Ex_\sigma^{\neg\emptyset}$) problems for EAAF. For any complexity class C, C-c (resp. C-h) means C-complete (resp. C-hard); an interval C-h, C' means C-hard and in C'.

	AAF						EAAF			
σ	Ver_σ	CA_σ	SA_σ	Ex_σ	$Ex_\sigma^{\neg\emptyset}$	DS_σ	VE_σ	EA_σ	Ex_σ	$Ex_\sigma^{\neg\emptyset}$
gr	P	P	P	trivial	P	trivial	P	P	trivial	P
co	P	NP-c	P	trivial	NP-c	coNP-c	Θ_2^p-h, Δ_2^P	Θ_2^p-h, Δ_2^P	trivial	NP-c
pr	coNP-c	NP-c	Π_2^P-c	trivial	NP-c	Π_2^P-c	Θ_2^p-h, Δ_3^P	Π_2^P-h, Δ_3^P	trivial	NP-c

remains the same as that for AAF if we focus on the grounded semantics or on the (non-empty) existence problem, we found that the complexity generally increases w.r.t. that of AAF in the other cases, that is for the verification and acceptance problems under complete or preferred semantics. This particularly hold if we compare the complexity of VE_σ and EA_σ for EAAF with that of Ver_σ and CA_σ for AAF with $\sigma \in \{\text{co}, \text{pr}\}$, as well as comparing the complexity of EA_{co} with SA_{co}. Finally, deciding acceptance in EAAF is harder than (resp. at least hard as) checking determinism in AAF, which in turn is harder than (resp. at least hard as) deciding skeptical acceptance in AAF for complete (resp. preferred) semantics.

6 RELATED WORK

Several proposals have been made to extend Dung's framework with the aim of better modeling the knowledge to be represented. The extensions include Bipolar AAF [48, 60], AAF with recursive attacks and supports [29, 30, 40], Dialectical framework [25], Abstract Reasoning Framework [11], AAF with preferences [10, 12, 46] and constraints [16, 31], as well extensions for representing uncertain information, e.g. incomplete AAF [21] and probabilistic AAF [1, 34, 42, 44].

Work on epistemic logic dates back to the early 1860s. Since then epistemic logic has played an important role also in computer science. This field is very active and important results are reported in a recent book surveying state-of-the-art research [58].

Epistemic Logic extends propositional logic by allowing to also express knowledge of the agents, also called subjective knowledge. In particular, it allows writing formule of the form $K_a\varphi$ stating that agent a knows that formula φ is true. Thus, while a formula $p \vee q$ expresses the (objective) knowledge that p or q are true (i.e. the formula is satisfied if a is true or b is true), a formula $K_a(p \vee q)$ expresses the (subjective) knowledge that agent a knows that p is true or q is true. On the other side, the formula K_aK_bp states that agent a knows that agent b knows that p holds, i.e. it is satisfied if agent a knows that agent b knows that p holds [59].

The idea of extending logic with epistemic constructs has been investigated also in the field of Answer Set Programming (ASP) [37–39]. Epistemic logic programs, firstly proposed in [38], extend disjunctive logic programs under the stable model semantics with modal constructs called subjective literals [26, 27, 37, 39]. The introduction of this extension was originally motivated to correctly represent incomplete information in programs that have several stable models. Using subjective literals, it is possible to check whether a literal is true in every or some stable model of the program. These models in this context are also called belief sets, being collected in a set called world view. The main idea was to expand the syntax and semantics of Answer Set Programming by modal operators **K** and **M** where $\mathbf{K}\varphi$ holds if φ is true in all answer sets of a program and $\mathbf{M}\varphi$ holds if φ is true in at least one answer set. Using this notation, *not* $\mathbf{K}p \wedge$ *not* $\mathbf{K} -p$ would correspond to "the truth value of p is unknown" even in the presence of multiple answer sets. In such a context, several problems are still open and they regard the support required by stable models, as well as splitting properties that are satisfied by classical ASP semantics, but not satisfied by epistemic ASP-based semantics [27, 41, 55].

As discussed in Section 2.2, in Epistemic Argumentation Framework (EAF) [54] the epistemic constraint plays a role similar to the one played by constraints in Constrained AAF (CAF) [6, 16, 31] or to the one played by preferences in Preference-based AAF (PAF) [8, 12, 15, 43], that is, restrict the set of feasible extensions. But, while in CAF and PAF the subset obtained is unique, in EAF we can have alternative subsets. It is worth noting that, according to our complexity analysis, EAAF cannot be fully captured in EAF. Indeed, the verification problem, that is, given an EAF and a set of arguments S, checking whether there is a complete epistemic labelling set containing S is in P for EAF. In contrast, as shown earlier, the verification problem for complete semantics is Θ_2^p-hard for EAAF. Similar differences arise for preferred semantics, entailing a different expressive power of the considered frameworks.

Although our focus is on argumentation, we believe that our results could be of interest to the logic community. In fact, by exploiting the correspondence between AF and Logic Programming [5], the proposed EAAF semantics could be seen as an alternative semantics for a special class of epistemic logic programs whose complexity and computation can be characterized by using our results.

7 CONCLUSION AND FUTURE WORK

We have presented the Epistemic Abstract Argumentation Framework, a generalization of Dung's framework where epistemic attacks and arguments can be expressed. We provided clear semantics for well-formed EAAF and an algorithm for its computation. We also provided complexity bounds for the verification, existence, and acceptance problems in EAAF under three well-known argumentation semantics. Our complexity analysis shows that the epistemic elements (i.e., epistemic attacks/arguments) impact on the complexity of most of the problems considered, and thus, it turns out that EAAF generally increases the expressivity of AAFs.

Future work will be devoted to considering other argumentation semantics such as the stable semantics. However, it is worth noting that the stable semantics for EAAF is intrinsically defined in the paper for odd-cycle free EAAF as, in this case, it collapses to the preferred semantics.

ACKNOWLEDGMENTS

We acknowledge financial support from PNRR MUR project PE0000013-FAIR.

REFERENCES

[1] G. Alfano, M. Calautti, S. Greco, F. Parisi, and I. Trubitsyna. 2020. Explainable Acceptance in Probabilistic Abstract Argumentation: Complexity and Approximation. In *Proc. of KR*. 33–43.
[2] Gianvincenzo Alfano, Andrea Cohen, Sebastian Gottifredi, Sergio Greco, Francesco Parisi, and Guillermo Ricardo Simari. 2020. Dynamics in Abstract Argumentation Frameworks with Recursive Attack and Support Relations. In *Proc. of ECAI*. 577–584.
[3] Gianvincenzo Alfano, Sergio Greco, and Francesco Parisi. 2021. Incremental Computation in Dynamic Argumentation Frameworks. *IEEE Intell. Syst.* 36, 6 (2021), 80–86.
[4] Gianvincenzo Alfano, Sergio Greco, Francesco Parisi, Gerardo I. Simari, and Guillermo Ricardo Simari. 2021. Incremental computation for structured argumentation over dynamic DeLP knowledge bases. *Artif. Intell.* 300 (2021), 103553.
[5] Gianvincenzo Alfano, Sergio Greco, Francesco Parisi, and Irina Trubitsyna. 2020. On the Semantics of Abstract Argumentation Frameworks: A Logic Programming Approach. *Theory Pract. Log. Program.* 20, 5 (2020), 703–718.
[6] Gianvincenzo Alfano, Sergio Greco, Francesco Parisi, and Irina Trubitsyna. 2021. Argumentation Frameworks with Strong and Weak Constraints: Semantics and Complexity. In *Proc. of AAAI*. 6175–6184.
[7] Gianvincenzo Alfano, Sergio Greco, Francesco Parisi, and Irina Trubitsyna. 2022. Incomplete Argumentation Frameworks: Properties and Complexity. In *Proc. of AAAI*. 5451–5460.
[8] Gianvincenzo Alfano, Sergio Greco, Francesco Parisi, and Irina Trubitsyna. 2022. On Preferences and Priority Rules in Abstract Argumentation. In *Proc. of IJCAI*. 2517–2524.
[9] Gianvincenzo Alfano, Sergio Greco, Francesco Parisi, and Irina Trubitsyna. 2023. Abstract Argumentation Framework with Conditional Preferences. In *Proc. of AAAI*. (to appear).
[10] Gianvincenzo Alfano, Sergio Greco, Francesco Parisi, and Irina Trubitsyna. 2023. Abstract Argumentation Framework with Conditional Preferences. In *Proc. of AAAI*.
[11] Gianvincenzo Alfano, Sergio Greco, Francesco Parisi, and Irina Trubitsyna. 2023. On acceptance conditions in abstract argumentation frameworks. *Information Sciences* 625 (2023), 757–779.
[12] Leila Amgoud and Claudette Cayrol. 2002. Inferring from Inconsistency in Preference-Based Argumentation Frameworks. *J. Autom. Reason.* 29, 2 (2002), 125–169.
[13] Leila Amgoud, Yannis Dimopoulos, and Pavlos Moraitis. 2007. A unified and general framework for argumentation-based negotiation. In *Proc. of AAMAS*. 158.
[14] Leila Amgoud and Mathieu Serrurier. 2008. Agents that argue and explain classifications. *Auton. Agents Multi Agent Syst.* 16, 2 (2008), 187–209.
[15] Leila Amgoud and Srdjan Vesic. 2014. Rich preference-based argumentation frameworks. *Int. J. Approx. Reason.* 55, 2 (2014), 585–606.
[16] Ofer Arieli. 2015. Conflict-free and conflict-tolerant semantics for constrained argumentation frameworks. *J. Appl. Log.* 13, 4 (2015), 582–604.
[17] K. Atkinson, P. Baroni, M. Giacomin, A. Hunter, H. Prakken, C. Reed, G. R. Simari, M. Thimm, and S. Villata. 2017. Towards Artificial Argumentation. *Artificial Intelligence Magazine* 38, 3 (2017), 25–36.
[18] Pietro Baroni and Massimiliano Giacomin. 2007. On principle-based evaluation of extension-based argumentation semantics. *Artif. Intell.* 171, 10-15 (2007), 675–700.
[19] P. Baroni, M. Giacomin, and G. Guida. 2005. SCC-recursiveness: a general schema for argumentation semantics. *Artificial Intelligence* 168, 1-2 (2005), 162–210.
[20] Pietro Baroni, Massimiliano Giacomin, and Beishui Liao. 2014. On topology-related properties of abstract argumentation semantics. A correction and extension to Dynamics of argumentation systems: A division-based method. *Artif. Intell.* 212 (2014), 104–115.
[21] D. Baumeister, M. Järvisalo, D. Neugebauer, A. Niskanen, and J. Rothe. 2021. Acceptance in incomplete argumentation frameworks. *Artif. Intell.* (2021), 103470.
[22] T.J.M. Bench-Capon and P. E. Dunne. 2007. Argumentation in Artificial Intelligence. *Artif. Intell.* 171 (2007), 619 – 641.
[23] Michael Bernreiter, Wolfgang Dvorák, and Stefan Woltran. 2022. Abstract Argumentation with Conditional Preferences. In *Proc. of COMMA*. 92–103.
[24] Stefano Bistarelli, Lars Kotthoff, Francesco Santini, and Carlo Taticchi. 2021. Summary Report for the Third International Competition On Computational Models of Argumentation. *AI Mag.* 42, 3 (2021), 70–73.
[25] G. Brewka, H. Strass, S. Ellmauthaler, J. P. Wallner, and S. Woltran. 2013. Abstract Dialectical Frameworks Revisited. In *Proc. of IJCAI*. 803–809.
[26] Pedro Cabalar, Jorge Fandinno, and Luis Fariñas del Cerro. 2020. Autoepistemic answer set programming. *Artif. Intell.* 289 (2020), 103382.
[27] Pedro Cabalar, Jorge Fandinno, and Luis Fariñas del Cerro. 2021. Splitting Epistemic Logic Programs. *Theory Pract. Log. Program.* 21, 3 (2021), 296–316.
[28] Álvaro Carrera and Carlos Angel Iglesias. 2015. A systematic review of argumentation techniques for multi-agent systems research. *Artif. Intell. Rev.* 44, 4 (2015), 509–535.
[29] C. Cayrol, J. Fandinno, L. F. del Cerro, and M.-C. Lagasquie-Schiex. 2018. Structure-Based Semantics of Argumentation Frameworks with Higher-Order Attacks and Supports. In *Proc. of COMMA*. 29–36.
[30] A. Cohen, S. Gottifredi, A. J. Garcia, and G. R. Simari. 2015. An approach to abstract argumentation with recursive attack and support. *J. Appl. Log.* 13, 4 (2015), 509–533.
[31] Sylvie Coste-Marquis, Caroline Devred, and Pierre Marquis. 2006. Constrained Argumentation Frameworks. In *Proc. of (KR)*. 112–122.
[32] Yannis Dimopoulos, Jean-Guy Mailly, and Pavlos Moraitis. 2019. Argumentation-based Negotiation with Incomplete Opponent Profiles. In *Proc. AAMAS*. 1252–1260.
[33] Phan Minh Dung. 1995. On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games. *Artif. Intell.* 77 (1995), 321–358.
[34] P. M. Dung and P. M. Thang. 2010. Towards (Probabilistic) Argumentation for Jury-based Dispute Resolution. In *Proc. of COMMA*. 171–182.
[35] Wolfgang Dvorák and Paul E. Dunne. 2017. Computational Problems in Formal Argumentation and their Complexity. *FLAP* 4, 8 (2017).
[36] Xiuyi Fan and Francesca Toni. 2012. Agent Strategies for ABA-based Information-seeking and Inquiry Dialogues. In *Proc. of ECAI*, Vol. 242. 324–329.
[37] Jorge Fandinno, Wolfgang Faber, and Michael Gelfond. 2021. Thirty years of Epistemic Specifications. *CoRR* abs/2108.07669 (2021).
[38] Michael Gelfond. 1991. Strong Introspection. In *Proc AAAI Conf.*, Thomas L. Dean and Kathleen R. McKeown (Eds.). 386–391.
[39] Michael Gelfond. 2011. New Semantics for Epistemic Specifications. In *Proc. of LPNMR Conf.*, Vol. 6645. 260–265.
[40] S. Gottifredi, A. Cohen, A. J. Garcia, and G. R. Simari. 2018. Characterizing acceptability semantics of argumentation frameworks with recursive attack and support relations. *Artif. Intell.* 262 (2018), 336–368.
[41] Andreas Herzig and Antonio Yuste-Ginel. 2021. On the Epistemic Logic of Incomplete Argumentation Frameworks. In *Proc. of KR*. 681–685.
[42] A. Hunter. 2012. Some Foundations for Probabilistic Abstract Argumentation. In *Proc. of COMMA*. 117–128.
[43] Souhila Kaci, Leendert W. N. van der Torre, and Serena Villata. 2018. Preference in Abstract Argumentation. In *Proc. COMMA*. 405–412.
[44] H. Li, N. Oren, and T. J. Norman. 2011. Probabilistic Argumentation Frameworks. In *Proc. of TAFA*. 1–16.
[45] Peter McBurney and Simon Parsons. 2009. Dialogue Games for Agent Argumentation. In *Argumentation in Artificial Intelligence*. Springer, 261–280.
[46] Sanjay Modgil and Henry Prakken. 2013. A general account of argumentation with preferences. *Artif. Intell.* 195 (2013), 361–397.
[47] Andreas Niskanen and Matti Järvisalo. 2020. μ-toksia: An Efficient Abstract Argumentation Reasoner. In *Proc. ofKR*. 800–804.
[48] F. Nouioua. 2013. AFs with Necessities: Further Semantics and Labelling Characterization. In *Proc. of SUM*. 120–133.
[49] Santiago Ontañón and Enric Plaza. 2007. An Argumentation-Based Framework for Deliberation in Multi-agent Systems. In *Int. Workshop, ArgMAS*, Vol. 4946. 178–196.
[50] Christos H. Papadimitriou. 1994. *Computational complexity.* Addison-Wesley.
[51] Simon Parsons, Carles Sierra, and Nicholas R. Jennings. 1998. Agents That Reason and Negotiate by Arguing. *J. Log. Comput.* 8, 3 (1998), 261–292.
[52] Henry Prakken. 2009. Models of Persuasion Dialogue. In *Argumentation in Artificial Intelligence*. 281–300.
[53] Henry Prakken and Giovanni Sartor. 1998. Modelling Reasoning with Precedents in a Formal Dialogue Game. *Artif. Intell. Law* 6, 2-4 (1998), 231–287.
[54] Chiaki Sakama and Tran Cao Son. 2020. Epistemic Argumentation Framework: Theory and Computation. *J. Artif. Intell. Res.* 69 (2020), 1103–1126.
[55] Yi-Dong Shen and Thomas Eiter. 2020. Constraint Monotonicity, Epistemic Splitting and Foundedness Are Too Strong in Answer Set Programming. *CoRR* abs/2010.00191 (2020).
[56] Yoav Shoham and Kevin Leyton-Brown. 2009. *Multiagent Systems - Algorithmic, Game-Theoretic, and Logical Foundations.* Cambridge University Press.
[57] G. R. Simari and I. Rahwan (Eds.). 2009. *Argumentation in Artificial Intelligence.* Springer.
[58] Hans van Ditmarsch, Joseph Y. Halpern, Wiebe van der Hoek, and Barteld P. Kooi (Eds.). 2015. *Handbook of Epistemic Logic.* kings College Pubn.
[59] Hans van Ditmarsch, Joseph Y. Halpern, Wiebe van der Hoek, and Barteld P. Kooi. 2015. An Introduction to Logics of Knowledge and Belief. *CoRR* abs/1503.00806 (2015). http://arxiv.org/abs/1503.00806
[60] S. Villata, G. Boella, D. M. Gabbay, and L. W. N. van der Torre. 2012. Modelling defeasible and prioritized support in bipolar argumentation. *Ann. Math. Artif. Intell.* 66, 1-4 (2012).