Abstract Argumentation Framework with Priority Rules and Preferences

Gianvincenzo Alfano, Sergio Greco, Francesco Parisi and Irina Trubitsyna

Department of Informatics, Modeling, Electronics and System Engineering (DIMES), University of Calabria, Rende, Italy

Abstract

Dung's abstract Argumentation Framework (AF) has emerged as a central formalism for argumentation in AI. In this paper, we discuss a recently proposed framework called *AF with Priority rules* (AFP) [1] that extends AF with sequences of *priority rules* which are able to express several kinds of desiderata among AF extensions. Using AFP, AF semantics can be viewed as ways to express priorities among extensions. We extend AFP by presenting *AF with Priority rules and Preferences* (AFP^2), where also preferences over arguments can be used to define priority rules. We study the complexity of the verification as well as credulous and skeptical acceptance problems for AFP and AFP^2.

Keywords

Abstract Argumentation, Priorities, Preferences, Computational Complexity.

1. Introduction

Abstract argumentation has emerged as one of the major fields in AI [2]. In particular, recent years have witnessed intensive formal study, development and application of Dung's abstract Argumentation Framework (AF) in various directions [3, 4, 5, 6, 7, 8, 9, 10, 11, 12]. Dung's framework is recognized as a simple, yet powerful formalism for modelling disputes between two or more agents. An AF consists of a set A of arguments and an attack relation $\Omega \subseteq A \times A$ that specifies conflicts between arguments (if argument a attacks argument b, then b is acceptable only if a is not). We can think of an AF as a directed graph whose nodes represent arguments and edges represent attacks. The meaning of an AF is given in terms of argumentation semantics, e.g. the well-known *grounded* ($\mathtt{gr}$), *complete* ($\mathtt{co}$) *preferred* ($\mathtt{pr}$), *stable* ($\mathtt{st}$), and *semi-stable* ($\mathtt{ss}$) semantics, which intuitively tell us the sets of arguments (called *σ-extensions*, with $\sigma \in \{\mathtt{gr}, \mathtt{co}, \mathtt{pr}, \mathtt{st}, \mathtt{ss}\}$) that can collectively be accepted to support a point of view in a dispute. For instance, for AF $\langle A, \Omega\rangle = \langle\{\mathtt{a}, \mathtt{b}\}, \{(\mathtt{a}, \mathtt{b}), (\mathtt{b}, \mathtt{a})\}\rangle$ having two arguments, $\mathtt{a}$ and $\mathtt{b}$, attacking each other, there are two preferred/stable extensions, $\{\mathtt{a}\}$ and $\{\mathtt{b}\}$, and neither argument $\mathtt{a}$ nor $\mathtt{b}$ is skeptically accepted. To cope with such situations, a possible solution is to provide means for preferring one argument to another, as shown in the following example.

Discussion Papers - 21st International Conference of the Italian Association for Artificial Intelligence (AIxIA 2022), November 28- December 2, 2022, Udine, Italy

g.alfano@dimes.unical.it (G. Alfano); greco@dimes.unical.it (S. Greco); fparisi@dimes.unical.it (F. Parisi); i.trubitsyna@dimes.unical.it (I. Trubitsyna)

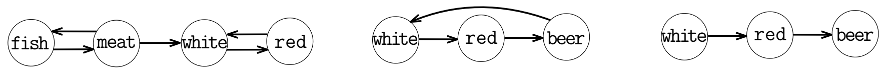

Figure 1: AF Λ_1 of Example 1 (left), AFs Λ_2 (center) and $\overline{\Lambda_2}$ (right) of Example 2.

Example 1. Consider the AF $\Lambda_1 = \langle\{\mathtt{fish}, \mathtt{meat}, \mathtt{red}, \mathtt{white}\}, \{(\mathtt{fish}, \mathtt{meat}), (\mathtt{meat}, \mathtt{fish}), (\mathtt{meat}, \mathtt{white}), (\mathtt{white}, \mathtt{red}), (\mathtt{red}, \mathtt{white})\}\rangle$, whose corresponding graph is shown in Figure 1(left-hand side). Intuitively, Λ_1 describes what a person is going to have for lunch. (S)he will have either `fish` or `meat`, and will drink either `white` wine or `red` wine. However, if (s)he will have `meat`, then (s)he will not drink `white` wine. Λ_1 has six complete extensions $E_0 = \emptyset$, $E_1 = \{\mathtt{fish}, \mathtt{white}\}$, $E_2 = \{\mathtt{fish}, \mathtt{red}\}$, $E_3 = \{\mathtt{meat}, \mathtt{red}\}$, $E_4 = \{\mathtt{fish}\}$, and $E_5 = \{\mathtt{red}\}$, which represent possible menus; E_0 is the grounded extension, whereas E_1, E_2 and E_3 are stable, preferred and semi-stable extensions. Assume now that person prefers to have `meat` instead of `fish` as main dish. Under such an assumption there is only one stable (and preferred) extension, namely E_3, which in a sense satisfies the person's preference. □

AF has been extended to Preference-based Argumentation Framework (PAF) where preferences stating that an argument is better than another are considered. Two main approaches have been proposed in the literature to define PAF semantics. The first approach defines PAF semantics in terms of that of an *auxiliary AF* [13, 14, 15]. However, there are cases where this semantics may give counterintuitive results as shown next.

Example 2. Consider a PAF consisting of the AF $\Lambda_2 = \langle\{\mathtt{white}, \mathtt{red}, \mathtt{beer}\}, \{(\mathtt{white}, \mathtt{red}), (\mathtt{red}, \mathtt{beer}), (\mathtt{beer}, \mathtt{white})\}\rangle$ shown in Figure 1(center), and the preference $\mathtt{white} > \mathtt{beer}$ that intuitively states that `white` is better than `beer`. According to the first approach for defining PAF semantics, for the auxiliary AF $\overline{\Lambda_2}$ shown in Figure 1(right-hand side), obtained from Λ_2 by removing attack $(\mathtt{beer}, \mathtt{white})$ conflicting with preference $\mathtt{white} > \mathtt{beer}$, there is only one complete extension, that is $\{\mathtt{white}, \mathtt{beer}\}$. However, this is not an extension for the underlying AF Λ_2 as it is not conflict-free w.r.t. Λ_2 (since `beer` attacks `white`). □

Herein, the problem is that preferences and attacks describe different pieces of knowledge and should be considered separately. This is carried out by the second approach for defining PAF semantics that compares extensions w.r.t. preferences defined over arguments [14, 15]. Following this approach, we introduce a general framework for dealing with preferences and priority rules in AF.

Contribution. We first discuss *AF with Priority rules* (AFP) [1] which extends AF with sequences of priority rules allowing to reasoning about extensions. We show that AFP generalizes AF with the classical semantics (i.e., $\mathtt{gr}, \mathtt{co}, \mathtt{pr}, \mathtt{st}, \mathtt{ss}$). Encoding such argumentation semantics in AFP means expressing priorities on the complete extensions of the underlying AF. Next, results concerning the complexity of the verification as well as the credulous and skeptical acceptance problems in AFP are given in Section 3.3.

Then, in Section 3.4, PAF and AFP are combined by extending AFP with preferences between arguments that lead to preferences between extensions (with the same spirit of PAF). The resulting framework, called *AF with Priority rules and Preferences* (AFP^2), is able to capture existing and novel PAF semantics. Finally, the complexity of the above-mentioned problems for the case of AFP^2 framework is studied. Notably, the complexity of AFP^2 does not increase w.r.t. that of AFP.

We assume the reader is familiar with the complexity classes used in the paper.

2. Preliminaries

We review the Dung's framework and its generalization with preferences (PAF).

2.1. Abstract Argumentation Framework

An abstract *Argumentation Framework* (AF) is a pair $\langle A, \Omega \rangle$, where A is a (finite) set of *arguments* and $\Omega \subseteq A \times A$ is a set of *attacks* (also called *defeats*). Different semantics have been defined for AF leading to the characterization of collectively acceptable sets of arguments, called *extensions* [16].

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

- $\mathit{Def}(E) = \{a \in A \mid \exists b \in E \,.\, (b, a) \in \Omega\}$;
- $\mathit{Acc}(E) = \{a \in A \mid \forall b \in A \,.\, (b, a) \in \Omega \Rightarrow b \in \mathit{Def}(E)\}$.
- $\mathit{Undec}(E) = A \setminus (E \cup \mathit{Def}(E))$.

To simplify the notation, we will often use E^+ and E^u to denote $\mathit{Def}(E)$ and $\mathit{Undec}(E)$, respectively.

Given an AF $\langle A, \Omega \rangle$, a set $E \subseteq A$ of arguments is said to be *conflict-free* iff $E \cap E^+ = \emptyset$; *admissible* iff it is conflict-free and $E \subseteq \mathit{Acc}(E)$. Given an AF $\langle A, \Omega \rangle$, a set $E \subseteq A$ is an *extension* called:

- *complete* ($\mathtt{co}$) iff it is conflict free and $E = \mathit{Acc}(E)$;
- *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 $E \cup E^+ = A$ or, equivalently, $E^u = \emptyset$;
- *semi-stable* ($\mathtt{ss}$) iff it is a complete extension with a minimal set of undecided arguments, i.e., E^u is $\subseteq$-minimal;
- *grounded* ($\mathtt{gr}$) iff it is the $\subseteq$-smallest complete extension.

The set of complete (resp. preferred, stable, semi-stable, grounded) extensions of an AF Λ will be denoted by $\mathtt{co}(\Lambda)$ (resp. $\mathtt{pr}(\Lambda)$, $\mathtt{st}(\Lambda)$, $\mathtt{ss}(\Lambda)$, $\mathtt{gr}(\Lambda)$). With a little abuse of notation, in the following we also use $\mathtt{gr}(\Lambda)$ to denote the grounded extension. 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 *multiple status* semantics. For any AF Λ, $\mathtt{st}(\Lambda) \subseteq \mathtt{ss}(\Lambda) \subseteq \mathtt{pr}(\Lambda) \subseteq \mathtt{co}(\Lambda)$ and $\mathtt{gr}(\Lambda) \in \mathtt{co}(\Lambda)$. Note that stable (resp.

semi-stable) extensions could be also defined as *preferred* extensions containing an empty (resp. minimal) set of undecided arguments.

Example 3. Let $\Lambda_3 = \langle A_3, \Omega_3 \rangle$ be an AF where $A_3 = \{\mathtt{a}, \mathtt{b}, \mathtt{c}, \mathtt{d}\}$ and $\Omega_3 = \{(\mathtt{a}, \mathtt{b}), (\mathtt{b}, \mathtt{a}), (\mathtt{a}, \mathtt{c}), (\mathtt{b}, \mathtt{c}), (\mathtt{c}, \mathtt{d}), (\mathtt{d}, \mathtt{c})\}$. The grounded extension is $\emptyset$ whereas the preferred extensions are $\{\mathtt{a}, \mathtt{d}\}$ and $\{\mathtt{b}, \mathtt{d}\}$, which are also stable and semi-stable. □

Given an AF $\Lambda = \langle A, \Omega \rangle$ and a semantics $\sigma \in \{\mathtt{co}, \mathtt{pr}, \mathtt{st}, \mathtt{ss}, \mathtt{gr}\}$, the *verification* problem, denoted as Ver_σ, is deciding whether a set $S \subseteq A$ is a σ-extension of Λ. Moreover, for a goal argument $g \in A$, the *credulous* (resp. *skeptical*) *acceptance* problem, denoted as CA_σ (resp. SA_σ), is deciding whether g belongs to *any* (resp. *every*) σ-extension of Λ. Clearly, $CA_{\mathtt{gr}}$ and $SA_{\mathtt{gr}}$ are identical problems.

2.2. Preference-based AFs

Several works generalizing Dung's AF to handle preferences over arguments have been proposed [17, 13, 18, 14, 19, 20].

Definition 1. *A* Preference-based Argumentation Framework (PAF) *is a triple* $\langle A,\Omega,>\rangle$ *such that* $\langle A, \Omega \rangle$ *is an AF and* $>$ *is a strict partial order (i.e. an irreflexive, asymmetric and transitive relation) over A, called preference relation.*

For arguments a and b, $a > b$ means that a is better than b. Two main approaches have been proposed to handle preferences in argumentation.

The first approach considers AF-based semantics and consists in first defining a defeat relation Ω_i that combines attacks in Ω and preference relations, and then applying the usual semantics on the AF $\langle A, \Omega_i \rangle$. Here Ω_i (with $i \in [1, 4]$) denotes one of the four mappings proposed in the literature [13, 14, 15]. As discussed in the Introduction, in some cases these semantics fail to capture the expected meaning and, therefore, we will not further discuss them. We point out that the complexity of acceptance problems does not increase as the mapping to AF (i.e., building Ω_i) is polynomial time.

The second approach to handle preferences considers extensions selection semantics for PAF [14, 15]. Here, given a PAF $\langle A, \Omega, > \rangle$, classical argumentation semantics are used to obtain extensions of the underlying AF $\langle A, \Omega \rangle$, and then the preference relation $>$ is used to obtain a preference relation $\succeq$ over such extensions, so that the *best* extensions w.r.t. $\succeq$ are eventually selected. There have been different proposals to define the best extensions, corresponding to different criteria to define $\succeq$.

Definition 2. *Given a PAF* $\langle A, \Omega, > \rangle$*, for* $E, F \subseteq A$ *with* $E \neq F$*, we have that under*

- democratic (d) *criterion [14]:* $E \succeq F$ *if* $\forall b \in F \setminus E\ \exists a \in E \setminus F$ *such that* $a > b$*;*
- elitist (e) *criterion [14]:* $E \succeq F$ *if* $\forall a \in E \setminus F\ \exists b \in F \setminus E$ *such that* $a > b$*;*
- KTV (k) *criterion [15]:* $E \succeq F$ *if* $\forall a, b \in A$ *the relation* $a > b$ *with* $a \in F \setminus E$ *and* $b \in E \setminus F$ *does not hold.*

Moreover, $E \succ F$*, if* $E \succeq F$ *and* $F \not\succeq E$*.*

Definition 3. *Given a PAF $\Delta = \langle A, \Omega, > \rangle$, a semantics $\sigma \in \{\mathtt{gr}, \mathtt{co}, \mathtt{pr}, \mathtt{st}, \mathtt{ss}\}$, and a criterion $* \in \{d, e, k\}$ for $\succeq$, the best σ-extensions of Δ under criterion $*$ (denoted as $\sigma_*(\Delta)$) are the extensions $E \in \sigma(\langle A, \Omega\rangle)$ such that there is no $F \in \sigma(\langle A, \Omega\rangle)$ with $F \succ E$.*

Example 4. Consider the following three PAFs: $\Delta_1 = \langle\{\mathtt{a}, \mathtt{b}, \mathtt{c}\}, \{(\mathtt{a}, \mathtt{b}), (\mathtt{b}, \mathtt{a}), (\mathtt{a}, \mathtt{c})\}, \{\mathtt{a} > \mathtt{b}\}\rangle$, $\Delta_2 = \langle\{\mathtt{a}, \mathtt{b}, \mathtt{d}\}, \{(\mathtt{a}, \mathtt{b}), (\mathtt{b}, \mathtt{a}), (\mathtt{b}, \mathtt{d})\}, \{\mathtt{a} > \mathtt{b}\}\rangle$, $\Delta_3 = \langle\{\mathtt{a}, \mathtt{b}, \mathtt{c}, \mathtt{d}\}, \{(\mathtt{a}, \mathtt{b}), (\mathtt{b}, \mathtt{a}), (\mathtt{a}, \mathtt{c}), (\mathtt{b}, \mathtt{d})\}, \{\mathtt{a} > \mathtt{b}\}\rangle$. The preferred extensions for the underlying AFs Λ_i, obtained from Δ_i by ignoring the preferences, are: $\mathtt{pr}(\Lambda_1) = \{E_1 = \{\mathtt{a}\}, E_2 = \{\mathtt{b}, \mathtt{c}\}\}$; $\mathtt{pr}(\Lambda_2) = \{E_3 = \{\mathtt{a}, \mathtt{d}\}, E_4 = \{\mathtt{b}\}\}$; $\mathtt{pr}(\Lambda_3) = \{E_3 = \{\mathtt{a}, \mathtt{d}\}, E_2 = \{\mathtt{b}, \mathtt{c}\}\}$.

The best preferred extensions are: $\mathtt{pr}_e(\Delta_1) = \mathtt{pr}_k(\Delta_1) = \{E_1\}$, $\mathtt{pr}_d(\Delta_1) = \{E_1, E_2\}$; $\mathtt{pr}_d(\Delta_2) = \mathtt{pr}_k(\Delta_2) = \{E_3\}$, $\mathtt{pr}_e(\Delta_2) = \{E_3, E_4\}$; $\mathtt{pr}_k(\Delta_3) = \{E_3\}$, $\mathtt{pr}_e(\Delta_3) = \mathtt{pr}_d(\Delta_3) = \{E_3, E_2\}$. □

An alternative definition for PAF, based on that defined in [21] for logic programs with preferences, has been proposed in [22]. In this context a PAF is a triple $\langle A, \Omega, \geq\rangle$, where $\geq$ is a reflexive and transitive relation and $a > b$ if $a \geq b$ and $b \not\geq a$. Moreover, the preference relation $\succeq$ over extensions is reflexive ($E \succeq E$), transitive ($E \succeq F$ and $F \succeq G$ implies $E \succeq G$) and satisfies the condition $E \succeq F$ if $\exists a \in E \setminus F, \exists b \in F \setminus E$ such that $a \geq b$ and $\nexists c \in F \setminus E$ such that $c > a$. In this paper we only deal with PAFs where relation $\succeq$ is not transitive as our proposal is intended to extend PAF, where $\succeq$ is not transitive for all the criteria of Definition 2 (e.g. KTV).

Observe that the preference relation makes sense only for multiple-status semantics, i.e. semantics prescribing more than one extension. In fact, for the unique-status grounded semantics, $\mathtt{gr}_*(\langle A, \Omega, >\rangle) = \mathtt{gr}(\langle A, \Omega\rangle)$ with $* \in \{d, e, k\}$.

Verification and Credulous/Skeptical Acceptance Problems. The verification problem for PAF, denoted as Ver_σ with $\sigma \in \{\mathtt{co}_*, \mathtt{pr}_*, \mathtt{st}_*, \mathtt{ss}_*, \mathtt{gr}_*\}$ and $* \in \{d, e, k\}$, extends that for AF by considering best extensions. Given a PAF $\Delta = \langle A, \Omega, >\rangle$, Ver_σ consists in checking whether a set $S \subseteq A$ belongs to $\sigma(\Delta)$, where $\sigma(\Delta)$ is the set of best σ-extensions of PAF Δ. Similarly, for a goal argument $g \in A$, the credulous (resp. skeptical) acceptance problem, denoted as CA_σ (resp. SA_σ), consists in deciding whether g belongs to any (resp. every) σ-extension in $\sigma(\Delta)$.

The complexity of the verification as well as credulous and skeptical acceptance problems for PAF under the *democratic*, *elitist*, and *KTV* criteria for multi-status semantics $\sigma \in \{\mathtt{co}, \mathtt{pr}, \mathtt{st}, \mathtt{ss}\}$ is presented in [1]. It turns out that the complexity of these problems generally increases of one level in the polynomial hierarchy w.r.t. the corresponding problems for AF.

3. AF with Priority Rules

In this section we extend AF with priority rules that allow us to express several kinds of desiderata among extensions, e.g. expressing classical AF semantics. Preferences between arguments are then considered in Subsection 3.4.

3.1. Syntax

A priority rule defines a priority between two extensions on the base of the satisfaction of a first order formula. Our formulae are built by considering variables denoting sets of arguments and variables denoting single arguments, logical connectives $\wedge, \vee$ and $\neg$, built-in predicates and functions operating on sets of arguments as described next.

The vocabulary consists of finite sets of (constant) arguments, argument variables, set variables, built-in predicates and functions and natural numbers in the interval $[0, |A|]$, where A is the set of arguments. In the following, arguments, argument variables, and set variables are denoted by lowercase letters a, b, c, d, lowercase letters x, y, z, and uppercase letters E, F, G, respectively. Therefore, we have simple terms (constant arguments and variable arguments) and set terms (set variables). The built-in (binary, infix) predicates are:

- $\in$ (predicate in): $x \in E$ checks if x belongs to set term E;
- comparison predicates $>, \geq, <, \leq$, to compare natural numbers (got by cardinality function applied to sets, see below);
- comparison predicates $=$ and $\neq$ to compare terms.

The built-in functions are *Acc*, *Def* and *Undec* defined earlier for AFs and the unary *cardinality function* $|S|$ computing the number of elements in S.

Definition 4. *For an AF $\Lambda = \langle A, \Omega \rangle$, a priority rule is an expression of the form $E \sqsupseteq F \leftarrow body$, where E and F are two distinct set variables and body is a quantified first order formula using simple terms, set variables E and F, predicates and functions, where E and F range over* $\texttt{co}(\Lambda)$*, and argument variables range over A.*

Example 5. Examples of priority rules are: $\varphi_1 = E \sqsupseteq F \leftarrow \forall x \,.\, \neg(x \in F) \vee (x \in E)$; $\varphi_2 = E \sqsupseteq F \leftarrow \forall x \,.\, \neg(x \in E^+) \vee (x \in F^+)$; and $\varphi_3 = E \sqsupseteq F \leftarrow |E| \geq |F|$. □

We use E^+ and E^u as shorthand for *Def*(E) and *Undec*(E). We also use the shorthand $\notin$ since $x \notin E \equiv \neg(x \in E)$. Finally, we may use the predicates $\subset, \subseteq$ to compare set terms as shorthands for the corresponding quantified first order formulae, e.g. $F \subseteq E \equiv \forall x \,.\, x \notin F \vee x \in E$.

Definition 5. *An* AF with Priority rules *(*AFP*) is a triple $\langle A, \Omega, \Phi \rangle$, where $\langle A, \Omega \rangle$ is an AF and $\Phi = [\varphi_1, ..., \varphi_n]$ is a linearly ordered set of priority rules (with $n \geq 0$).*

3.2. Semantics

The semantics of AFPs is given by extensions which are 'prioritized' w.r.t. partially ground instances of priority rules, as explained in what follows.

For any AFP $\Delta = \langle A, \Omega, \Phi \rangle$, let $\Lambda = \langle A, \Omega \rangle$ be the AF associated with Δ, $pground_\Lambda(\Phi)$ (or simply $pground(\Phi)$ whenever Λ is understood) denotes the set of partially grounded priority rules derived from Φ by replacing head set variables with constant set terms (i.e. complete extensions). Furthermore, $ground_\Lambda(\Phi)$ denotes the set of ground rules derived from $pground_\Lambda(\Phi)$ by making variable-free the body of priority rules, as illustrated in the following example.

Example 6. Consider the AFP $\Delta_6 = \langle A_6, \Omega_6, \Phi_6 \rangle$ with $A_6 = \{\texttt{a}, \texttt{b}, \texttt{c}\}$, $\Omega_6 = \{(\texttt{a}, \texttt{b}), (\texttt{b}, \texttt{a})\}$, and $\Phi_6 = [E \sqsupseteq F \leftarrow \forall x.(x \notin F) \vee (x \in E)]\rangle$. Here, set variables E and F take values from $\texttt{co}(\langle A_6, \Omega_6 \rangle) = \{\{\texttt{c}\}, \{\texttt{a}, \texttt{c}\}, \{\texttt{b}, \texttt{c}\}\}$. For the partially grounded priority rule:

$\{\mathtt{a},\mathtt{c}\} \sqsupseteq \{\mathtt{c}\} \leftarrow \forall \mathtt{x}.(\mathtt{x} \notin \{\mathtt{c}\}) \vee (\mathtt{x} \in \{\mathtt{a},\mathtt{c}\})$, the ground rule is as follows:

$\{\mathtt{a},\mathtt{c}\} \sqsupseteq \{\mathtt{c}\} \leftarrow ((\mathtt{a} \notin \{\mathtt{c}\}) \vee (\mathtt{a} \in \{\mathtt{a},\mathtt{c}\})) \wedge ((\mathtt{b} \notin \{\mathtt{c}\}) \vee (\mathtt{b} \in \{\mathtt{a},\mathtt{c}\})) \wedge$
$((\mathtt{c} \notin \{\mathtt{c}\}) \vee (\mathtt{c} \in \{\mathtt{a},\mathtt{c}\}))$.

The body of the ground rule is true. Its intuitive meaning is that $\{\mathtt{a},\mathtt{c}\}$ is "better" than $\{\mathtt{c}\}$. □

Before defining the semantics of an AFP, we introduce some notations. Let $\langle A, \Omega, [\varphi]\rangle$ be an AFP, $\mathcal{C} = \mathtt{co}(\langle A, \Omega\rangle)$, and $E, F \in \mathcal{C}$ two complete extensions. Then $E \succeq F$ w.r.t. φ if there exists a partially ground instantiation of φ of the form $E \sqsupseteq F \leftarrow body$ such that $body$ evaluates to true. Moreover, $E \succ F$ (w.r.t. φ) if $E \succeq F$ and $F \not\succeq E$; $E \in \mathcal{C}$ is a *prioritized extension* w.r.t. φ if there exists no extension $F \in \mathcal{C}$ such that $F \succ E$. We use $\beta_\varphi(\mathcal{C})$ to denote the set of prioritized extensions in $\mathcal{C}$ w.r.t. φ.

Definition 6. *Given an AFP $\Delta = \langle A, \Omega, \Phi = [\varphi_1, ..., \varphi_n]\rangle$, the set of prioritized extensions of Δ w.r.t. Φ is given by $\beta_{\varphi_n}(...\beta_{\varphi_1}(\mathtt{co}(\langle A, \Omega\rangle)...)$ and is denoted by $\mathtt{co}(\langle A, \Omega, \Phi\rangle)$.*

We do not consider transitivity of relation $\sqsupseteq$ and focus on explicit prioritized rules stating e.g. E is as good as F. A transitive closure of $\sqsupseteq$ would require to (iteratively) adding a ground prioritized rule $E \sqsupseteq F \leftarrow body_1, body_2$ for each pair of ground rules $E \sqsupseteq G \leftarrow body_1$ and $G \sqsupseteq F \leftarrow body_2$, which can yield an exponential blow-up in the number of prioritized rules. Nonetheless, if needed, transitivity can still be stated by including the transitive closure in Φ.

Encoding AF semantics in AFP. As shown below, AF semantics can be easily expressed in AFP; the encoding for st, that may admit no extensions, is given separately. As shown in [1], AFP also encodes several cardinality-based semantics for AF.

Proposition 1. *For AF $\Lambda = \langle A, \Omega\rangle$ and $\sigma \in \{\mathtt{gr}, \mathtt{pr}, \mathtt{ss}\}$, it holds that $\sigma(\Lambda) = \mathtt{co}(\langle A, \Omega, [\varphi_\sigma]\rangle)$ with: $\varphi_{\mathtt{gr}} = E \sqsupseteq F \leftarrow F \supseteq E$; $\varphi_{\mathtt{pr}} = E \sqsupseteq F \leftarrow F \subseteq E$; $\varphi_{\mathtt{ss}} = E \sqsupseteq F \leftarrow E^u \subseteq F^u$.*

Proposition 2. *For any AF $\Lambda = \langle A, \Omega\rangle$, let $A' = A \cup \{\alpha, \overline{\alpha}\}$ and $\Omega' = \Omega \cup \{(\alpha, \overline{\alpha}), (\overline{\alpha}, \alpha)\} \cup \{(\alpha, a) \mid a \in A\}$. Let $\varphi_{\mathtt{st}} = E \sqsupseteq F \leftarrow E^u \subseteq F^u \wedge \overline{\alpha} \in E$. It holds that $\mathtt{st}(\Lambda) = \emptyset$ if $\{\alpha\} \in \mathtt{co}(\langle A', \Omega', [\varphi_{\mathtt{st}}]\rangle)$; otherwise $\mathtt{st}(\Lambda) = \{E \setminus \{\overline{\alpha}\} \mid E \in \mathtt{co}(\langle A', \Omega', [\varphi_{\mathtt{st}}]\rangle)\}$.*

3.3. Acceptance and Verification Problems in AFP

Given an AFP Δ and a set S of arguments, the *prioritized verification* problem, denoted as PV, is the problem of deciding whether $S \in \mathtt{co}(\Delta)$, i.e. S is a prioritized extension of Δ. Moreover, given an argument g, the *prioritized credulous (resp. skeptical) acceptance* problem, denoted as PCA (resp. PSA), is the problem of deciding whether g belongs to any (resp. all) prioritized extension in $\mathtt{co}(\Delta)$.

Theorem 1. *For any AFP $\langle A, \Omega, \Phi\rangle$, PV (resp. PCA, PSA) is in $\Pi^p_{|\Phi|}$ (resp. $\Sigma^p_{|\Phi|+1}$, $\Pi^p_{|\Phi|+1}$).*

In our complexity analysis the input consists of three sets and its size is $|A| + |\Omega| + |\Phi|$. That is, the number of variables in the body of a rule is considered bounded by a constant, i.e. not part of the input, thus grounding a rule as well as its evaluation is polynomial. Though this can be seen as a limitation, in practice, the number of variables needed in a rule can be reasonably bounded by a constant. As a matter of fact, at most two variables per rule are used in all our examples and in the semantics encodings in Propositions 1 and 2, as well as in Proposition 3 below.

Tighter complexity bounds can be obtained by using the result of Proposition 1 that entails that for any AFP $\langle A, \Omega, \Phi \rangle$, with $|\Phi| = 1$, PV (resp. PCA, PSA) is coNP-complete (resp. Σ_2^p-complete, Π_2^p-complete). Specifically, the hardness results can be shown by providing a reduction from $Ver_{\texttt{ss}}$ (resp. $CA_{\texttt{ss}}$, $SA_{\texttt{ss}}$) for AF [23] to our problem with $\Phi = [\varphi_{\texttt{ss}}]$.

3.4. Combining Preferences with Priorities

We extend AFP with preferences between arguments. Specifically, we allow the use of the predicate $>$ introduced for PAF to compare arguments in the body of priority rules.

Definition 7. *An* AF with Priority rules and Preferences *(AFP^2) is a tuple $\langle A, \Omega, \Phi, > \rangle$, where $\langle A, \Omega, \Phi \rangle$ is an AFP and $>$ is a strict partial order over A.*

Example 7. The priority rule $E \sqsupseteq F \leftarrow \exists x, y \,.\, (x \in E) \wedge (y \in F) \wedge (x > y)$, which uses preferences among arguments, states that E is as good as F if there is an argument in E which is preferred to an argument in F. □

The following proposition states that PAF semantics can be encoded in AFP^2. Particularly, the set of best σ-extensions of a given PAF can be defined by filtering out from the set of complete extension of an AFP^2 those that follow the priority rules (*i*) φ_σ encoding the chosen complete-based semantics σ (cf. Proposition 1), and (*ii*) φ_* encoding one of the preference criteria (i.e. deterministic, elitist and KTV of Definition 2).

Proposition 3. *For any PAF $\Delta = \langle A, \Omega, > \rangle$, $* \in \{d, e, k\}$ and $\sigma \in \{\texttt{co}, \texttt{gr}, \texttt{pr}, \texttt{ss}\}$, it holds that $\sigma_*(\Delta) = \texttt{co}(\langle A, \Omega, [\varphi_\sigma, \varphi_*], > \rangle)$ where φ_σ is empty for $\sigma = \texttt{co}$ and as defined in Proposition 1 for $\sigma \in \{\texttt{gr}, \texttt{pr}, \texttt{ss}\}$, and:*
- $\varphi_d = E \sqsupseteq F \leftarrow \forall y \in F \setminus E \;\; \exists x \in E \setminus F \,.\, x > y;$
- $\varphi_e = E \sqsupseteq F \leftarrow \forall x \in E \setminus F \;\; \exists y \in F \setminus E \,.\, x > y;$
- $\varphi_k = E \sqsupseteq F \leftarrow \neg(\exists x \in (F \setminus E) \, \exists y \in (E \setminus F) \,.\, x > y).$

Moreover, $\texttt{st}_*(\Delta) = \emptyset$ if $\{\alpha\} \in \texttt{co}(\langle A', \Omega', [\varphi_{\texttt{st}}] \rangle)$; otherwise $\texttt{st}_*(\Delta) = \{E \setminus \{\overline{\alpha}\} \mid E \in \texttt{co}(\langle A', \Omega', [\varphi_{\texttt{st}}, \varphi_*] \rangle)\}$, where $A', \Omega', \varphi_{\texttt{st}}$ are as in Proposition 2.

Interestingly, the complexity of AFP^2 does not increase w.r.t. that of AFP [1]. We have that, for any AFP^2 $\langle A, \Omega, \Phi, > \rangle$, *PV* (resp. *PCA*, *PSA*) is in $\Pi^p_{|\Phi|}$ (resp. in $\Sigma^p_{|\Phi|+1}$, in $\Pi^p_{|\Phi|+1}$).

We believe that the idea behind AFP^2 concerning priorities on extensions, i.e. preferences between solutions, could be explored for structured argumentation formalisms [24, 25, 26, 27, 28, 29, 30] where preferences are typically used to resolve attacks into defeats between arguments. As for implementations of our framework, given the connection between AF semantics and LP models [31, 32], ASP systems such as *DLV* and *potassco* that support cardinality-based semantics can be used to define encodings for AFP semantics by extending those for AF [33]. Finally, we plan to investigate preferences (possibly conditioned ones [34, 35, 36]) in incomplete AF [10, 37], probabilistic AF [38, 39, 40, 41], and AF with constraints [9].

References

[1] G. Alfano, S. Greco, F. Parisi, I. Trubitsyna, On preferences and priority rules in abstract argumentation, in: Proc. of IJCAI, 2022, pp. 2517–2524.

[2] D. Gabbay, M. Giacomin, G. R. Simari, M. Thimm (Eds.), Handbook of Formal Argumentation, volume 2, College Public., 2021.

[3] D. Baumeister, M. Järvisalo, D. Neugebauer, A. Niskanen, J. Rothe, Acceptance in incomplete argumentation frameworks, Artif. Intell. (2021) 103470.

[4] G. Alfano, S. Greco, F. Parisi, A meta-argumentation approach for the efficient computation of stable and preferred extensions in dynamic bipolar argumentation frameworks, Intelligenza Artificiale 12 (2018) 193–211.

[5] G. Alfano, S. Greco, F. Parisi, On scaling the enumeration of the preferred extensions of abstract argumentation frameworks, in: Proc. of ACM SAC, 2019, pp. 1147–1153.

[6] G. Alfano, A. Cohen, S. Gottifredi, S. Greco, F. Parisi, G. R. Simari, Dynamics in abstract argumentation frameworks with recursive attack and support relations, in: Proc. of ECAI, 2020, pp. 577–584.

[7] G. Alfano, S. Greco, F. Parisi, Incremental computation in dynamic argumentation frameworks, IEEE Intell. Syst. 36 (2021) 80–86.

[8] G. Alfano, S. Greco, Incremental skeptical preferred acceptance in dynamic argumentation frameworks, IEEE Intell. Syst. 36 (2021) 6–12.

[9] G. Alfano, S. Greco, F. Parisi, I. Trubitsyna, Argumentation frameworks with strong and weak constraints: Semantics and complexity, in: Proc. of AAAI, 2021, pp. 6175–6184.

[10] B. Fazzinga, S. Flesca, F. Furfaro, Revisiting the notion of extension over incomplete abstract argumentation frameworks, in: Proc. of IJCAI, 2020, pp. 1712–1718.

[11] B. Fazzinga, S. Flesca, F. Furfaro, Abstract argumentation frameworks with marginal probabilities, in: Proc. of IJCAI, 2022, pp. 2613–2619.

[12] B. Fazzinga, S. Flesca, F. Furfaro, L. Pontieri, Process mining meets argumentation: Explainable interpretations of low-level event logs via abstract argumentation, Inf. Syst. 107 (2022) 101987.

[13] L. Amgoud, C. Cayrol, Inferring from inconsistency in preference-based argumentation frameworks, J. Autom. Reason. 29 (2002) 125–169.

[14] L. Amgoud, S. Vesic, Rich preference-based argumentation frameworks, Int. J. Approx. Reason. 55 (2014) 585–606.

[15] S. Kaci, L. W. N. van der Torre, S. Villata, Preference in abstract argumentation, in: Proc. COMMA, 2018, pp. 405–412.

[16] P. M. Dung, On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games, Artif. Intell. 77 (1995) 321–358.

[17] L. Amgoud, C. Cayrol, On the acceptability of arguments in preference-based argumentation, in: Proc. of UAI, 1998, pp. 1–7.

[18] L. Amgoud, S. Vesic, A new approach for preference-based argumentation fra-meworks, Ann. Math. Artif. Intell. 63 (2011) 149–183.

[19] K. Cyras, Argumentation-based reasoning with preferences, in: PAAMS, 2016, pp. 199–210.

[20] R. Silva, S. Sá, J. F. L. Alcântara, Semantics hierarchy in preference-based argumentation

frameworks, in: COMMA, 2020, pp. 339–346.
[21] C. Sakama, K. Inoue, Prioritized logic programming and its application to commonsense reasoning, Artif. Intell. 123 (2000).
[22] T. Wakaki, Preference-based argumentation built from prioritized logic programming, J. Log. Comput. 25 (2015) 251–301.
[23] W. Dvorák, P. E. Dunne, Computational problems in formal argumentation and their complexity, FLAP 4 (2017).
[24] H. Prakken, G. Sartor, Argument-based extended logic programming with defeasible priorities, J. Appl. Non Class. Logics 7 (1997) 25–75.
[25] S. Modgil, H. Prakken, A general account of argumentation with preferences, Artif. Intell. 195 (2013) 361–397.
[26] F. Toni, A tutorial on assumption-based argumentation, Arg. & Comp. 5 (2014) 89–117.
[27] G. Alfano, S. Greco, F. Parisi, G. I. Simari, G. R. Simari, Incremental computation of warranted arguments in dynamic defeasible argumentation: the rule addition case, in: Proc. of ACM SAC, 2018, pp. 911–917.
[28] A. J. Garcia, H. Prakken, G. R. Simari, A comparative study of some central notions of ASPIC+ and DeLP, Theory Pract. Log. Program. 20 (2020) 358–390.
[29] J. Heyninck, C. Strasser, A comparative study of assumption-based argumentative approaches to reasoning with priorities, FLAP 8 (2021) 737–808.
[30] G. Alfano, S. Greco, F. Parisi, G. I. Simari, G. R. Simari, Incremental computation for structured argumentation over dynamic DeLP knowledge bases, AI 300 (2021) 103553.
[31] M. Caminada, S. Sá, J. F. L. Alcântara, W. Dvorák, On the equivalence between logic programming semantics and argumentation semantics, IJAR 58 (2015) 87–111.
[32] G. Alfano, S. Greco, F. Parisi, I. Trubitsyna, On the semantics of abstract argumentation frameworks: A logic programming approach, TPLP 20 (2020) 703–718.
[33] W. Dvorák, S. A. Gaggl, A. Rapberger, J. P. Wallner, S. Woltran, The ASPARTIX system suite, in: Proc. of COMMA, 2020, pp. 461–462.
[34] G. Alfano, S. Greco, F. Parisi, I. Trubitsyna, Abstract argumentation framework with conditional preferences, in: Proc. of AAAI, 2023, p. (to appear).
[35] T. Lukasiewicz, E. Malizia, Complexity results for preference aggregation over (m)cp-nets: Pareto and majority voting, Artif. Intell. 272 (2019) 101–142.
[36] T. Lukasiewicz, E. Malizia, Complexity results for preference aggregation over (m)cp-nets: Max and rank voting, Artif. Intell. 303 (2022) 103636.
[37] G. Alfano, S. Greco, F. Parisi, I. Trubitsyna, Incomplete argumentation frameworks: Properties and complexity, in: Proc. of AAAI, 2022, pp. 5451–5460.
[38] B. Fazzinga, S. Flesca, F. Parisi, On the complexity of probabilistic abstract argumentation frameworks, ACM Trans. Comput. Log. 16 (2015) 22:1–22:39.
[39] B. Fazzinga, S. Flesca, F. Furfaro, Complexity of fundamental problems in probabilistic abstract argumentation: Beyond independence, Artif. Intell. 268 (2019) 1–29.
[40] B. Fazzinga, S. Flesca, F. Furfaro, Probabilistic bipolar abstract argumentation frameworks: complexity results, in: Proc. of IJCAI, 2018, pp. 1803–1809.
[41] G. Alfano, M. Calautti, S. Greco, F. Parisi, I. Trubitsyna, Explainable acceptance in probabilistic abstract argumentation: Complexity and approximation, in: Proc. of KR, 2020, pp. 33–43.