On Preferences and Priority Rules in Abstract Argumentation

Gianvincenzo Alfano, **Sergio Greco**, **Francesco Parisi** and **Irina Trubitsyna**
DIMES Department, University of Calabria, Rende, Italy
{g.alfano, greco, fparisi, i.trubitsyna}@dimes.unical.it

Abstract

Dung's abstract Argumentation Framework (AF) has emerged as a central formalism for argumentation in AI. Preferences in AF allow to represent the comparative strength of arguments in a simple yet expressive way. In this paper we first investigate the complexity of the verification as well as credulous and skeptical acceptance problems in Preference-based AF (PAF) that extends AF with preferences over arguments. Next, after introducing new semantics for AF where extensions are selected using cardinality (instead of set inclusion) criteria and investigating their complexity, we introduce a framework called *AF with Priority rules* (AFP) that extends AF with sequences of *priority rules*. AFP generalizes AF with classical set-inclusion and cardinality based semantics, suggesting that argumentation semantics can be viewed as ways to express priorities among extensions. Finally, we extend AFP by proposing *AF with Priority rules and Preferences* (AFP^2), where also preferences over arguments can be used to define priority rules, and study the complexity of the above-mentioned problems.

1 Introduction

Recent years have witnessed intensive formal study, development and application of Dung's abstract Argumentation Framework (AF) in various directions [Gabbay *et al.*, 2021]. 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* (gr), *complete* (co) *preferred* (pr), *stable* (st), and *semi-stable* (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, a and b, attacking each other, there are two preferred/stable extensions, {a} and {b}, and neither argument a nor 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.

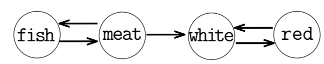

Figure 1: AF Λ of Example 1.

Example 1 Consider AF Λ=$\langle\{\mathtt{fish}, \mathtt{meat}, \mathtt{red}, \mathtt{white}\}, \{(\mathtt{fish}, \mathtt{meat}), (\mathtt{meat}, \mathtt{fish}), (\mathtt{meat}, \mathtt{white}), (\mathtt{white}, \mathtt{red}), (\mathtt{red}, \mathtt{white})\}\rangle$, whose corresponding graph is shown in Figure 1. Intuitively, Λ 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 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. $\square$

Two main approaches have been proposed in the literature for Preference-based Argumentation Frameworks (PAFs).

The first approach defines the PAF semantics in terms of an underlying AF [Amgoud and Cayrol, 2002a; Amgoud and Vesic, 2014; Kaci *et al.*, 2018]. Considering for instance the PAF of our previous example, where the preference is $\mathtt{meat} > \mathtt{fish}$ (meat is better than fish), the underlying AF according to the approach proposed in [Amgoud and Cayrol, 2002a] consists of an AF where the attack $(\mathtt{fish}, \mathtt{meat})$ is deleted. Consequently, there is only one complete extension (namely E_3) and both meat and red are skeptically accepted.

However, there are cases where this PAF semantics may give counterintuitive results as shown next.

Example 2 Consider a scenario where David wants to organize a party and would invite Anne, Bob and Cayrol, and he received the following answers: Bob will attend the party if Anne does not, Cayrol will attend if Bob does not, and Anne will attend if Cayrol does not attend. Regarding the contrast between Anne and Cayrol, David would prefer that Anne attends the party. This could be represented by the

PAF $\Delta = \langle \{\mathtt{a}, \mathtt{b}, \mathtt{c}\}, \{(\mathtt{a}, \mathtt{b}), (\mathtt{b}, \mathtt{c}), (\mathtt{c}, \mathtt{a})\}, \{\mathtt{a} > \mathtt{c}\}\rangle$, where argument $\mathtt{x}$ states that "*(the person whose names' initial is)* $\mathtt{x}$ *joins the party*" and $\mathtt{a} > \mathtt{c}$ encodes the preference. According to the previous PAF semantics, for any semantics $\sigma \in \{\mathtt{gr}, \mathtt{co}, \mathtt{pr}, \mathtt{st}, \mathtt{ss}\}$, there is only one σ-extension $\{\mathtt{a}, \mathtt{c}\}$ stating that is Anne and Cayrol will attend the party. □

Observe that, in Example 2, the resulting set $\{\mathtt{a}, \mathtt{c}\}$ does not give the desired intuition as Anne and Cayrol should not attend the party together in any case (because of attack $(\mathtt{c}, \mathtt{a})$ the resulting set is not conflict-free). 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 defining the PAF semantics, that consists in comparing extensions w.r.t. preferences obtained from those between arguments [Amgoud and Vesic, 2014; Kaci *et al.*, 2018].

In this paper we focus on the second PAF semantics, based on extension selection by means of preferences, and investigate AF with preferences from several standpoints including the complexity analysis of well-known reasoning problems as well as the introduction of a general framework for dealing with preferences in AF.

Contributions. Our main contributions are as follows.

- We study the complexity of the verification (Ver_σ) as well as credulous (CA_σ) and skeptical acceptance (SA_σ) problems for PAF under the *democratic* (d), *elitist* (e), and *KTV* (k) criteria for multi-status semantics $\sigma \in \{\mathtt{co}, \mathtt{st}, \mathtt{pr}, \mathtt{ss}\}$ (Section 3). As shown in Table 1 that summarizes our results, 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.
- In Section 4 we introduce cardinality based semantics for AF, namely *cardinality preferred* ($\mathtt{c\text{-}pr}$) and *cardinality semi-stable* ($\mathtt{c\text{-}ss}$), that refine the set-inclusion based semantics $\mathtt{pr}$ and $\mathtt{ss}$ by allowing to filter out extensions on the basis of the cardinality of the set of arguments belonging to an extension. Cardinality based semantics are meant to be complementary to set based semantics, providing the user with further preferential criteria to select among complete extensions. For these semantics, as well as for another cardinality based semantics that is shown to collapse to one of them, we investigate the relationships with classical AF semantics as well as the complexity of Ver_σ, CA_σ and SA_σ with $\sigma \in \{\mathtt{c\text{-}pr}, \mathtt{c\text{-}ss}\}$. Interestingly while the complexity of Ver_σ remains coNP-complete, as for the corresponding set-inclusion based semantics, CA_σ and SA_σ become easier (in Θ_2^P and Δ_2^P instead of Σ_2^P and Π_2^P, respectively) under the cardinality based semantics $\mathtt{c\text{-}pr}$ and $\mathtt{c\text{-}ss}$.
- We extend AF with sequences of priority rules allowing to reasoning about extensions (Section 5). We call the resulting framework *AF with Priority rules* (AFP). We show that AFP generalizes AF with the above discussed cardinality based semantics ($\mathtt{c\text{-}pr}$ and $\mathtt{c\text{-}ss}$) as well as with the classical semantics ($\mathtt{gr}, \mathtt{co}, \mathtt{pr}, \mathtt{st}, \mathtt{ss}$). That is, encoding such argumentation semantics in AFP means expressing priorities on the complete extensions of the underlying AF. Results concerning the complexity of the verification as well as the credulous and skeptical acceptance problems in AFP are given in Section 5.3. Finally, in Section 5.4, PAF and AFP are combined by extending AFP with preferences between arguments that lead to preferences between extensions (with the same spirit of PAF). We show that the resulting framework, called *AF with Priority rules and Preferences* (AFP^2), is able to capture existing and novel PAF semantics. The complexity of AFP^2 is also investigated.

2 Preliminaries

We review the Dung's framework and its generalization with preferences (PAF), and recall some basic complexity classes.

2.1 Abstract Argumentation Framework

An abstract *Argumentation Framework* (AF) 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*). An AF can be seen as a directed graph, whose nodes represent arguments and edges represent attacks. Different semantics have been defined for AF leading to the characterization of collectively acceptable sets of arguments, called *extensions* [Dung, 1995].

Given an AF $\Lambda = \langle \mathrm{A}, \Omega\rangle$ and a set $E \subseteq \mathrm{A}$ of arguments, an argument $a \in \mathrm{A}$ 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 $\forall b \in \mathrm{A}$ 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):[1]

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

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 $\langle \mathrm{A}, \Omega\rangle$, a set $E \subseteq \mathrm{A}$ 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 $\langle \mathrm{A}, \Omega\rangle$, a set $E \subseteq \mathrm{A}$ 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 s.t. $E \cup E^+ = \mathrm{A}$ 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*

[1]For $E \in \mathtt{co}(\Lambda)$, $\mathit{Acc}(E)$ (resp. $\mathit{Def}(E)$, $\mathit{Undec}(E)$) is the set of arguments labelled as *in* (resp. *out*, *undec*) [Caminada, 2006].

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 $\Lambda = \langle \mathrm{A}, \Omega \rangle$ be an AF where $\mathrm{A} = \{\mathtt{a}, \mathtt{b}, \mathtt{c}, \mathtt{d}\}$ and $\Omega = \{(\mathtt{a}, \mathtt{b}), (\mathtt{b}, \mathtt{a}), (\mathtt{a}, \mathtt{c}), (\mathtt{b}, \mathtt{c}), (\mathtt{c}, \mathtt{d}), (\mathtt{d}, \mathtt{c})\}$. 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. $\square$

Given an AF $\Lambda = \langle \mathrm{A}, \Omega \rangle$ 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 $S \subseteq \mathrm{A}$ 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 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 [Amgoud and Cayrol, 1998; 2002a; Amgoud and Vesic, 2011; 2014; Cyras, 2016; Silva *et al.*, 2020].

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.*[2]

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 $\langle \mathrm{A}, \Omega_i \rangle$. Here Ω_i (with $i \in [1, 4]$) denotes one of the four mappings proposed in the literature [Amgoud and Cayrol, 2002a; Amgoud and Vesic, 2014; Kaci *et al.*, 2018]. 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 [Amgoud and Vesic, 2014; Kaci *et al.*, 2018]. 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 [Amgoud and Vesic, 2014]:* $E \succeq F$ *if* $\forall b \in F \setminus E\ \exists a \in E \setminus F$ *such that* $a > b$*;*
- elitist (e) *criterion [Amgoud and Vesic, 2014]:* $E \succeq F$ *if* $\forall a \in E \setminus F\ \exists b \in F \setminus E$ *such that* $a > b$*;*
- KTV (k) *criterion [Kaci* et al.*, 2018]:* $E \succeq F$ *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 as follows:
- $\mathtt{pr}_e(\Delta_1) = \mathtt{pr}_k(\Delta_1) = \{E_1\},\ \ \mathtt{pr}_d(\Delta_1) = \{E_1, E_2\}$;
- $\mathtt{pr}_d(\Delta_2) = \mathtt{pr}_k(\Delta_2) = \{E_3\},\ \ \mathtt{pr}_e(\Delta_2) = \{E_3, E_4\}$;
- $\mathtt{pr}_k(\Delta_3) = \{E_3\},\ \ \mathtt{pr}_e(\Delta_3) = \mathtt{pr}_d(\Delta_3) = \{E_3, E_2\}$. $\square$

An alternative definition for PAF, based on that defined in [Sakama and Inoue, 2000] for logic programs with preferences, has been proposed in [Wakaki, 2015]. In this context a PAF is a triple $\langle \mathrm{A}, \Omega, \geq \rangle$, 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 $\Delta = \langle \mathrm{A}, \Omega, > \rangle$, Ver_σ consists in checking whether a set $S \subseteq \mathrm{A}$ belongs to $\sigma(\Delta)$. Similarly, for a goal argument $g \in \mathrm{A}$, 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 $\sigma(\Delta)$.

2.3 Complexity Classes

We recall the main complexity classes used in the paper and, in particular, the definition of the classes Σ_h^P, Π_h^P and Δ_h^P, with $h \geq 0$ (see e.g. [Papadimitriou, 1994]):

- $\Sigma_0^P = \Pi_0^P = \Delta_0^P = P$; $\Sigma_1^P = NP$ 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

[2] Equivalent definitions use as a primitive the partial preorder $\geq$ and then derive $>$ [Amgoud and Vesic, 2014].

σ	Ver_σ	CA_σ	SA_σ
$\mathtt{co}_d$	coNP-c	Σ_2^P-c	Π_2^P-c
$\mathtt{co}_e$	*P*	*P*	*P*
$\mathtt{co}_k$	coNP-c	Σ_2^P-c	*P*
$\mathtt{st}_d, \mathtt{st}_e, \mathtt{st}_k$	coNP-c	Σ_2^P-c	Π_2^P-c
$\mathtt{pr}_d$	coNP-c	Σ_2^P-c	Π_2^P-c
$\mathtt{pr}_e, \mathtt{pr}_k$	Π_2^P-c	Σ_2^P-h, Σ_3^P	Π_2^P-h, Π_3^P
$\mathtt{ss}_d, \mathtt{ss}_e, \mathtt{ss}_k$	Π_2^P-c	Σ_2^P-h, Σ_3^P	Π_2^P-h, Π_3^P

Table 1: Complexity of the verification (Ver_σ), credulous (CA_σ) and skeptical (SA_σ) acceptance problems for PAF. 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'.

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 containing the problems that can be solved in polynomial time by a deterministic Turing machine by performing a number of calls bounded by $O(log\ n)$ to an oracle in the class Σ_{h-1}^P. It holds that $\Sigma_h^P \subseteq \Theta_{h+1}^P \subseteq \Delta_{h+1}^P \subseteq \Sigma_{h+1}^P \subseteq PSPACE$ and $\Pi_h^P \subseteq \Theta_{h+1}^P \subseteq \Delta_{h+1}^P \subseteq \Pi_{h+1}^P \subseteq PSPACE$.

3 Complexity of Acceptance Problems in PAF

We start with the following proposition showing that, irrespective of the preference relation $>$, best complete and grounded semantics for PAF coincide under elitist criterion, whereas best complete and best preferred semantics coincide under the democratic criterion; moreover, the grounded extension of the underlying AF is contained in the set of best complete extensions under KTV criterion.

Proposition 1 *For any PAF* $\Delta = \langle A, \Omega, > \rangle$ *it holds that* $i)$ $\mathtt{co}_e(\Delta) = \mathtt{gr}(\langle A, \Omega\rangle)$, $ii)$ $\mathtt{co}_d(\Delta) = \mathtt{pr}_d(\Delta)$, *and* $iii)$ $\mathtt{gr}(\langle A, \Omega\rangle) \in \mathtt{co}_k(\Delta)$.

We now characterize the complexity of the verification, credulous and skeptical acceptance problems for PAF. First, observe that the results of Proposition 1, and the fact that the grounded semantics prescribes exactly one extension, entail that $Ver_{\mathtt{co}_e}$, $CA_{\mathtt{co}_e}$, and $SA_{\mathtt{co}_e}$ are polynomial as these problems are equivalent to those for AF under the grounded semantics. Moreover, since $\mathtt{co}_d(\Delta) = \mathtt{pr}_d(\Delta)$ we will state the result for $\mathtt{pr}_d(\Delta)$ only.

Theorem 1 *Ver_σ is coNP-complete for* $\sigma \in \{\mathtt{co}_k, \mathtt{st}_d, \mathtt{st}_e, \mathtt{st}_k, \mathtt{pr}_d\}$, *and* Π_2^P*-complete for* $\sigma \in \{\mathtt{pr}_e, \mathtt{pr}_k, \mathtt{ss}_d, \mathtt{ss}_e, \mathtt{ss}_k\}$.

The complexity of the verification problem helps in addressing the complexity of the credulous and skeptical acceptance problems, which is given in the following theorems.

Theorem 2 *CA_σ is i)* Σ_2^P*-complete for* $\sigma \in \{\mathtt{co}_k, \mathtt{st}_d, \mathtt{st}_e, \mathtt{st}_k, \mathtt{pr}_d\}$*; ii)* Σ_2^P*-hard and in* Σ_3^P *for* $\sigma \in \{\mathtt{pr}_e, \mathtt{pr}_k, \mathtt{ss}_d, \mathtt{ss}_e, \mathtt{ss}_k\}$.

Theorem 3 *SA_σ is i) in P for* $\sigma \in \{\mathtt{co}_k\}$*; ii)* Π_2^P*-complete for* $\sigma \in \{\mathtt{st}_d, \mathtt{st}_e, \mathtt{st}_k, \mathtt{pr}_d\}$*; and iii)* Π_2^P*-hard and in* Π_3^P *for* $\sigma \in \{\mathtt{pr}_e, \mathtt{pr}_k, \mathtt{ss}_d, \mathtt{ss}_e, \mathtt{ss}_k\}$.

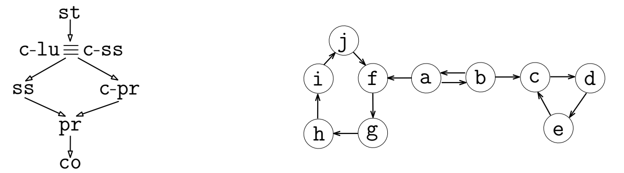

Figure 2: (Left) Relationships between cardinality based semantics and classical AF semantics: an arrow from σ_1 to σ_2 means that $\sigma_1(\Lambda) \subseteq \sigma_2(\Lambda)$. (Right) AF of Example 5.

The results of the previous theorems, summarized in Table 1, show that the complexity of the three problems generally increases of one level in the polynomial hierarchy w.r.t that of AFs for multiple-status semantics (see [Dvorák and Dunne, 2017] for a survey on the complexity of AF).

4 Cardinality Based Semantics for AF

In this section we present new semantics for AF that are inspired by the criteria at the basis of preferred and semi-stable semantics, which refine complete semantics by selecting from the set of complete extensions, respectively, the $\subseteq$-maximal sets and the sets having a $\subseteq$-minimal set of undecided arguments. In particular, instead of comparing sets by set inclusion, we focus on a filtering mechanism based on the cardinality of complete extensions. As it will be clear in what follows, our proposals for AF semantics refine the well-known preferred and semi-stable semantics.

Before providing the formal definitions, we introduce a partial order on pairs of (natural) numbers. Given the numbers a, b, c, d, we say that $i)$ $(a, b) \geq (c, d)$ iff $a \geq c$ and $b \geq d$, and $ii)$ $(a, b) > (c, d)$ iff $(a, b) \geq (c, d)$ and either $a > c$ or $b > d$. [3] Moreover, given two complete extensions E and F, we say that $E \succeq F$ (resp. $E \succ F$) if $(|E|, |E^+|) \geq (|F|, |F^+|)$ (resp. $(|E|, |E^+|) > (|F|, |F^+|)$).

Definition 4 *Given an AF* Λ*, a complete extension* $E \in \mathtt{co}(\Lambda)$ *is said to be:*

- cardinality preferred *(*$\mathtt{c\text{-}pr}$*) iff there is no complete extension* F *of* Λ *such that* $F \succ E$*;*
- cardinality semi-stable *(*$\mathtt{c\text{-}ss}$*) iff there is no complete extension* F *of* Λ *such that* $|F^u| < |E^u|$*;*
- cardinality least-undecided *(*$\mathtt{c\text{-}lu}$*) iff it is a cardinality preferred extension with a minimum number of undecided elements, i.e., there is no cardinality preferred extension* F *of* Λ *such that* $|F^u| < |E^u|$.

Therefore, (i) $\mathtt{c\text{-}pr}$ semantics returns the complete extensions whose numbers of accepted and defeated arguments are pairwise maximum (instead of $\subseteq$-maximal accepted arguments as for $\mathtt{pr}$); (ii) $\mathtt{c\text{-}ss}$ returns the complete extensions whose number of undecided elements is minimum (instead of $\subseteq$-minimal as for $\mathtt{ss}$); and (iii) $\mathtt{c\text{-}lu}$ refines $\mathtt{c\text{-}pr}$ by minimizing the number of undecided arguments.

[3] With a little abuse of notation, we use the same symbols $\leq$ and $<$ to compare numbers and pairs of numbers, though $\leq$ defines a linear order on numbers and a partial order on pairs of numbers.

Cardinality based semantics are useful when the user would prefer "optimal" solutions. As as example, consider the AF obtained by adding argument $\mathtt{sorbet}$ and attack $(\mathtt{meat}, \mathtt{sorbet})$ to the AF of Example 1. There are three stable, preferred and semi-stable extensions, $E_1 = \{\mathtt{fish}, \mathtt{white}, \mathtt{sorbet}\}$, $E_2 = \{\mathtt{fish}, \mathtt{red}, \mathtt{sorbet}\}$ and $E_3 = \{\mathtt{meat}, \mathtt{red}\}$, representing possible menus. Assume now that the customer would like to have as many menu items as possible. Hence the best extensions according to the customer's criteria are E_1 and E_2. This is exactly what is obtained by using the cardinality preferred semantics (c-pr)—classical AF semantics are not able to express such kind of criteria.

The following proposition states the relationships between cardinality and classical semantics, summarized in Figure 2.

Proposition 2 *For any AF Λ, it holds that:*

$i)$ $\mathtt{c\text{-}lu}(\Lambda) = \mathtt{c\text{-}ss}(\Lambda)$*;*
$ii)$ $\mathtt{st}(\Lambda) \subseteq \mathtt{c\text{-}ss}(\Lambda) \subseteq \mathtt{ss}(\Lambda)$*;*
$iii)$ $\mathtt{c\text{-}ss}(\Lambda) \subseteq \mathtt{c\text{-}pr}(\Lambda) \subseteq \mathtt{pr}(\Lambda)$*;*
$iv)$ $\mathtt{st}(\Lambda) \neq \emptyset$ *iff* $\mathtt{st}(\Lambda) = \mathtt{c\text{-}ss}(\Lambda) = \mathtt{ss}(\Lambda)$.

As c-lu and c-ss semantics coincide, hereafter we do not explicitly deal with c-lu semantics.

Example 5 Continuing with Example 2, assume now that the party planner would like to invite 10 people (corresponding to arguments $\mathtt{a}, \mathtt{b}, .., \mathtt{j}$) and that the conditions for attending the party are represented by the AF shown in Figure 2. Then, there are three complete extensions $E_1 = \emptyset$, $E_2 = \{\mathtt{a}, \mathtt{g}, \mathtt{i}\}$ and $E_3 = \{\mathtt{b}, \mathtt{d}\}$, with $E_1^+ = \emptyset$, $E_2^+ = \{\mathtt{f}, \mathtt{h}, \mathtt{j}\}$ and $E_3^+ = \{\mathtt{c}, \mathtt{e}\}$. E_2 and E_3 are preferred and semi-stable extensions, but as $E_3 \succ E_2$ only E_3 is a c-pr and c-ss extension. □

For any AF Λ $\mathtt{ss}(\Lambda) \subseteq \mathtt{pr}(\Lambda)$, and Proposition 2 tells us that $\mathtt{c\text{-}pr}(\Lambda) \subseteq \mathtt{pr}(\Lambda)$. The next proposition shows that the sets $\mathtt{ss}(\Lambda)$, $\mathtt{c\text{-}pr}(\Lambda)$ and $\mathtt{pr}(\Lambda)$ are distinct.

Proposition 3

- *There exists an AF Λ s.t.* $\mathtt{ss}(\Lambda) \not\subseteq \mathtt{c\text{-}pr}(\Lambda)$.
- *There exists an AF Λ s.t.* $\mathtt{ss}(\Lambda) \not\supseteq \mathtt{c\text{-}pr}(\Lambda)$.
- *There exists an AF Λ s.t.* $(\mathtt{ss}(\Lambda) \cup \mathtt{c\text{-}pr}(\Lambda)) \subset \mathtt{pr}(\Lambda)$.

Complexity of cardinality based semantics. The next theorem states the complexity of the verification and credulous and skeptical acceptance problems for AF under the two novel cardinality based semantics, i.e. c-pr and c-ss.

Theorem 4

- *Ver_σ is coNP-complete for $\sigma \in \{\mathtt{c\text{-}pr}, \mathtt{c\text{-}ss}\}$.*
- *Both $CA_{\mathtt{c\text{-}pr}}$ and $SA_{\mathtt{c\text{-}pr}}$ are Θ_2^P-hard and in Δ_2^P.*
- *Both $CA_{\mathtt{c\text{-}ss}}$ and $SA_{\mathtt{c\text{-}ss}}$ are Θ_2^P-complete.*

It turns out that, under standard complexity assumptions, computing credulous and skeptical acceptance under the cardinality based semantics c-pr and c-ss is easier than under classical maximal-set semantics pr and ss.

5 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 the above-discussed cardinality based semantics as well as the classical AF semantics. Preferences between arguments are then considered in Subsection 5.4.

5.1 Syntax

Priority rules define 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 5 *For an AF $\Lambda = \langle A, \Omega \rangle$, a priority rule is 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* $\mathtt{co}(\Lambda)$*, and argument variables range over A.*

Example 6 Some examples of priority rules are:

$$\varphi_1 : \; E \sqsupseteq F \leftarrow \forall x \,.\, \neg(x \in F) \vee (x \in E)$$
$$\varphi_2 : \; E \sqsupseteq F \leftarrow \forall x \,.\, \neg(x \in E^+) \vee (x \in F^+)$$
$$\varphi_3 : \; E \sqsupseteq F \leftarrow |E| \geq |F|.$$
□

Recall that 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 6 *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 $n \geq 0$).*

5.2 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 $\Delta = \langle \mathrm{A}, \Omega, \Phi \rangle$, let $\Lambda = \langle \mathrm{A}, \Omega \rangle$ 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 7 Consider the AFP $\Delta = \langle \mathrm{A} = \{\mathtt{a},\mathtt{b},\mathtt{c}\},\ \Omega = \{(\mathtt{a},\mathtt{b}),(\mathtt{b},\mathtt{a})\},\ \Phi = [E \sqsupseteq F \leftarrow \forall x.(x \notin F) \vee (x \in E)]\rangle$. Here, set variables E and F take values from $\mathtt{co}(\langle \mathrm{A},\Omega\rangle) = \{\{\mathtt{c}\},\{\mathtt{a},\mathtt{c}\},\{\mathtt{b},\mathtt{c}\}\}$. 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:

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

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

Before defining the semantics of an AFP, we introduce some notations. Let $\langle \mathrm{A},\Omega,[\varphi]\rangle$ be an AFP, $\mathcal{C} = \mathtt{co}(\langle \mathrm{A},\Omega\rangle)$, 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 7 *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 explicitly including the transitive closure in Φ.

AF semantics can be easily expressed in AFP; the encoding for $\mathtt{st}$, that may admit no extensions, is given separately.

Proposition 4 *For any AF $\Lambda = \langle A,\Omega\rangle$ and $\sigma \in \{\mathtt{gr},\mathtt{pr},\mathtt{ss},\mathtt{c\text{-}pr},\mathtt{c\text{-}ss}\}$, it holds that $\sigma(\Lambda) = \mathtt{co}(\langle A,\Omega,[\varphi_\sigma]\rangle)$ with:*

- $\varphi_{\mathtt{gr}} = E \sqsupseteq F \leftarrow F \supseteq E;$
- $\varphi_{\mathtt{pr}} = E \sqsupseteq F \leftarrow F \subseteq E;$
- $\varphi_{\mathtt{ss}} = E \sqsupseteq F \leftarrow E^u \subseteq F^u;$
- $\varphi_{\mathtt{c\text{-}pr}} = E \sqsupseteq F \leftarrow (|F| > |E| \wedge |F^+| \geq |E^+|) \vee (|F| \geq |E| \wedge |F^+| > |E^+|);$
- $\varphi_{\mathtt{c\text{-}ss}} = E \sqsupseteq F \leftarrow |E^u| \leq |F^u|.$

Proposition 5 *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*

$$\mathtt{st}(\Lambda) = \begin{cases}\emptyset, & \text{if } \{\alpha\} \in \mathtt{co}(\langle A',\Omega',[\varphi_{\mathtt{st}}]\rangle); \\ \{E \setminus \{\overline{\alpha}\} \mid E \in \mathtt{co}(\langle A',\Omega',[\varphi_{\mathtt{st}}]\rangle)\}, & \text{otherwise.}\end{cases}$$

5.3 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 5 *For any AFP $\langle A,\Omega,\Phi\rangle$, PV (resp. PCA, PSA) is in $\Pi^P_{|\Phi|}$ (resp. $\Sigma^P_{|\Phi|+1}$, $\Pi^P_{|\Phi|+1}$).*

It is worth noting that in our complexity analysis the input consists of three sets and its size is $|\mathrm{A}| + |\Omega| + |\Phi|$. 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 4 and 5, as well as in Proposition 6 below.

Tighter complexity bounds can be obtained by using the result of Proposition 4 that entails that for any AFP $\langle \mathrm{A},\Omega,\Phi\rangle$, with $|\Phi| = 1$, PV (resp. PCA, PSA) is coNP-complete (resp. Σ^P_2-complete, Π^P_2-complete). Specifically, the hardness results can be shown by providing a reduction from $Ver_{\mathtt{ss}}$ (resp. $CA_{\mathtt{ss}}$, $SA_{\mathtt{ss}}$) for AF [Dvorák and Dunne, 2017] to our problem with $\Phi = [\varphi_{\mathtt{ss}}]$.

5.4 Combining Preferences with Priorities

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

Definition 8 *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 8 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 4), and (*ii*) φ_* encoding one of the preference criteria (i.e. deterministic, elitist and KTV of Definition 2).

Proposition 6 *For any PAF $\Delta = \langle A,\Omega,>\rangle$, $* \in \{d,e,k\}$ and $\sigma \in \{\mathtt{co},\mathtt{gr},\mathtt{pr},\mathtt{ss}\}$, it holds that $\sigma_*(\Delta) = \mathtt{co}(\langle A,\Omega,[\varphi_\sigma,\varphi_*],>\rangle)$ where φ_σ is empty for $\sigma = \mathtt{co}$ and as defined in Proposition 4 for $\sigma \in \{\mathtt{gr},\mathtt{pr},\mathtt{ss}\}$, 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, $\mathtt{st}_*(\Delta) = \emptyset$ if $\{\alpha\} \in \mathtt{co}(\langle A',\Omega',[\varphi_{\mathtt{st}}]\rangle)$; otherwise $\mathtt{st}_*(\Delta) = \{E \setminus \{\overline{\alpha}\} \mid E \in \mathtt{co}(\langle A',\Omega',[\varphi_{\mathtt{st}},\varphi_*]\rangle)\}$, where $A',\Omega',\varphi_{\mathtt{st}}$ are as in Proposition 5.

The complexity of AFP^2 does not increase w.r.t. that of AFP.

Proposition 7 *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 $\Pi^P_{|\Phi|+1}$).*

6 Conclusions and Future Works

We introduced a general argumentation framework (AFP^2) where priorities between complete extensions as well as preferences between individual arguments can be expressed.

AFP^2 extends the research generalizing the Dung's framework for allowing preferences over arguments [Amgoud and Cayrol, 2002a; 2002b; Kaci *et al.*, 2018; Modgil, 2009; Amgoud and Vesic, 2011]. Priorities and preferences in AFP^2 are handled as in logic programming, where in answer set and circumscription approaches they are exploited to filter out the best models [Brewka and Eiter, 1999; Delgrande *et al.*, 2003; Eiter *et al.*, 2003; Sakama and Inoue, 2000; Brewka *et al.*, 2003; Greco *et al.*, 2007] and in [Brewka *et al.*, 2015] users may also extend the ASP system to define their own preference ordering through a set of primitives. Preferences can be also expressed in value-based AFs [Bench-Capon, 2003; Dunne and Bench-Capon, 2004], where each argument is associated with a numeric value, and a set of possible orders (preferences) among the values is defined. In [Dunne *et al.*, 2011; Coste-Marquis *et al.*, 2012] weights are associated with attacks, and new semantics extending the classical ones are proposed—we plan to investigate the connection between the above-mentioned frameworks and AFP^2 in future work.

As for cardinality based semantics, cardinality-minimum sets of arguments have been considered in [Caminada and Dunne, 2020; Dvorák and Wallner, 2020] where the complexity of identifying strongly admissible sets with bounded or minimal size is investigated, and computational approaches are provided to minimize the steps of a discussion [Caminada *et al.*, 2016]. Similarly to our language, YALLA [de Saint-Cyr *et al.*, 2016] allows to write logic formulae useful for reasoning on extensions under AF semantics. However, while AFP^2 is designed to compare complete extensions, YALLA is designed for more general purposes, such as describing AF graphs, structural semantics (i.e. semantics based on the attack relation), and also changes in argumentation systems.

We are interested in exploring the relationship between AFP^2 and *rich PAF* [Amgoud and Vesic, 2014], an extension of PAF providing two separate sets of preferences: one containing preferences used to generate a "repaired" AF, and one containing preferences used to select the best extensions. Moreover, following [Mailly and Rossit, 2020] that recently investigates to what extent preferences can be blended with ranking semantics, it would be interesting to explore the relationship between AFP^2 and ranking semantics [Bonzon *et al.*, 2016]. Also, we believe that the idea behind AFP^2 concerning priorities on extensions, i.e. preferences between solutions, could be explored for structured argumentation formalisms [Prakken and Sartor, 1997; Modgil and Prakken, 2013; Toni, 2014; Garcia *et al.*, 2020; Heyninck and Strasser, 2021; Alfano *et al.*, 2021a] 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 [Caminada *et al.*, 2015; Alfano *et al.*, 2020b], ASP implementations such as *DLV* and *potassco* that support cardinality-based semantics can be used to define encodings for AFP semantics by extending those for AF [Dvorák *et al.*, 2020]. Finally, we plan to investigate preferences in incomplete AF [Fazzinga *et al.*, 2020; Alfano *et al.*, 2022], probabilistic AF [Fazzinga *et al.*, 2015; Alfano *et al.*, 2020a], and AF with constraints [Alfano *et al.*, 2021b].

References

[Alfano *et al.*, 2020a] G. Alfano, M. Calautti, S. Greco, F. Parisi, and I. Trubitsyna. Explainable acceptance in probabilistic abstract argumentation: Complexity and approximation. In *Proc. of KR*, pages 33–43, 2020.

[Alfano *et al.*, 2020b] Gianvincenzo Alfano, Sergio Greco, Francesco Parisi, and Irina Trubitsyna. On the semantics of abstract argumentation frameworks: A logic programming approach. *TPLP*, 20(5):703–718, 2020.

[Alfano *et al.*, 2021a] Gianvincenzo Alfano, Sergio Greco, Francesco Parisi, Gerardo I. Simari, and Guillermo Ricardo Simari. Incremental computation for structured argumentation over dynamic DeLP knowledge bases. *Artif. Intell.*, 300:103553, 2021.

[Alfano *et al.*, 2021b] Gianvincenzo Alfano, Sergio Greco, Francesco Parisi, and Irina Trubitsyna. Argumentation frameworks with strong and weak constraints: Semantics and complexity. In *Proc. of AAAI*, pages 6175–6184, 2021.

[Alfano *et al.*, 2022] Gianvincenzo Alfano, Sergio Greco, Francesco Parisi, and Irina Trubitsyna. Incomplete argumentation frameworks: Properties and complexity. *Preprint of AAAI*, https://tinyurl.com/ycyt7bas, 2022.

[Amgoud and Cayrol, 1998] L. Amgoud and C. Cayrol. On the acceptability of arguments in preference-based argumentation. In *Proc. of UAI*, pages 1–7, 1998.

[Amgoud and Cayrol, 2002a] Leila Amgoud and Claudette Cayrol. Inferring from inconsistency in preference-based argumentation frameworks. *J. Autom. Reason.*, 29(2):125–169, 2002.

[Amgoud and Cayrol, 2002b] Leila Amgoud and Claudette Cayrol. A reasoning model based on the production of acceptable arguments. *AMAI*, 34(1-3):197–215, 2002.

[Amgoud and Vesic, 2011] Leila Amgoud and Srdjan Vesic. A new approach for preference-based argumentation frameworks. *Ann. Math. Artif. Intell.*, 63(2):149–183, 2011.

[Amgoud and Vesic, 2014] Leila Amgoud and Srdjan Vesic. Rich preference-based argumentation frameworks. *Int. J. Approx. Reason.*, 55(2):585–606, 2014.

[Bench-Capon, 2003] Trevor J. M. Bench-Capon. Persuasion in practical argument using value-based argumentation frameworks. *J. Log. Comput.*, 13(3):429–448, 2003.

[Bonzon *et al.*, 2016] Elise Bonzon, Jérôme Delobelle, Sébastien Konieczny, and Nicolas Maudet. A comparative study of ranking-based semantics for abstract argumentation. In *Proc. of AAAI*, pages 914–920, 2016.

[Brewka and Eiter, 1999] Gerhard Brewka and Thomas Eiter. Preferred answer sets for extended logic programs. *Artif. Intell.*, 109(1-2):297–356, 1999.

[Brewka *et al.*, 2003] Gerhard Brewka, Ilkka Niemelä, and Miroslaw Truszczynski. Answer set optimization. In *Proc. IJCAI*, pages 867–872, 2003.

[Brewka *et al.*, 2015] Gerhard Brewka, James P. Delgrande, Javier Romero, and Torsten Schaub. asprin: Customizing answer set preferences without a headache. In *Proc. AAAI*, pages 1467–1474, 2015.

[Caminada and Dunne, 2020] Martin Caminada and Paul E. Dunne. Minimal strong admissibility: A complexity analysis. In *Proc. of COMMA*, pages 135–146, 2020.

[Caminada *et al.*, 2015] Martin Caminada, Samy Sá, João F. L. Alcântara, and Wolfgang Dvorák. On the equivalence between logic programming semantics and argumentation semantics. *Int. J. Approx. Reason.*, 58:87–111, 2015.

[Caminada *et al.*, 2016] Martin W. A. Caminada, Wolfgang Dvorák, and Srdjan Vesic. Preferred semantics as socratic discussion. *J. Log. Comput.*, 26(4):1257–1292, 2016.

[Caminada, 2006] Martin Caminada. On the issue of reinstatement in argumentation. In *JELIA*, pages 111–123, 2006.

[Coste-Marquis *et al.*, 2012] Sylvie Coste-Marquis, Sébastien Konieczny, Pierre Marquis, and Mohand Akli Ouali. Weighted attacks in argumentation frameworks. In *Proc. KR*, 2012.

[Cyras, 2016] Kristijonas Cyras. Argumentation-based reasoning with preferences. In *PAAMS*, pages 199–210, 2016.

[de Saint-Cyr *et al.*, 2016] Florence Dupin de Saint-Cyr, Pierre Bisquert, Claudette Cayrol, and Marie-Christine Lagasquie-Schiex. Argumentation update in YALLA (yet another logic language for argumentation). *Int. J. Approx. Reason.*, 75:57–92, 2016.

[Delgrande *et al.*, 2003] James P. Delgrande, Torsten Schaub, and Hans Tompits. A framework for compiling preferences in logic programs. *TPLP*, 3(2):129–187, 2003.

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

[Dunne and Bench-Capon, 2004] Paul E. Dunne and Trevor J. M. Bench-Capon. Complexity in value-based argument systems. In *Proc. of JELIA*, pages 360–371, 2004.

[Dunne *et al.*, 2011] Paul E. Dunne, Anthony Hunter, Peter McBurney, Simon Parsons, and Michael J. Wooldridge. Weighted argument systems: Basic definitions, algorithms, and complexity results. *Artif. Intell.*, 175(2):457–486, 2011.

[Dvorák and Dunne, 2017] Wolfgang Dvorák and Paul E. Dunne. Computational problems in formal argumentation and their complexity. *FLAP*, 4(8), 2017.

[Dvorák and Wallner, 2020] Wolfgang Dvorák and Johannes Peter Wallner. Computing strongly admissible sets. In *Proc. of COMMA*, pages 179–190, 2020.

[Dvorák *et al.*, 2020] Wolfgang Dvorák, Sarah Alice Gaggl, Anna Rapberger, Johannes Peter Wallner, and Stefan Woltran. The ASPARTIX system suite. In *Proc. of COMMA*, pages 461–462, 2020.

[Eiter *et al.*, 2003] Thomas Eiter, Wolfgang Faber, Nicola Leone, and Gerald Pfeifer. Computing preferred answer sets by meta-interpretation in answer set programming. *Theory Pract. Log. Program.*, 3(4-5):463–498, 2003.

[Fazzinga *et al.*, 2015] Bettina Fazzinga, Sergio Flesca, and Francesco Parisi. On the complexity of probabilistic abstract argumentation frameworks. *ACM Trans. Comput. Log.*, 16(3):22:1–22:39, 2015.

[Fazzinga *et al.*, 2020] Bettina Fazzinga, Sergio Flesca, and Filippo Furfaro. Revisiting the notion of extension over incomplete abstract argumentation frameworks. In *Proc. of IJCAI*, pages 1712–1718, 2020.

[Gabbay *et al.*, 2021] Dov Gabbay, Massimiliano Giacomin, Guillermo R. Simari, and Matthias Thimm, editors. *Handbook of Formal Argumentation*. College Public., 2021.

[Garcia *et al.*, 2020] Alejandro Javier Garcia, Henry Prakken, and Guillermo Ricardo Simari. A comparative study of some central notions of ASPIC+ and DeLP. *Theory Pract. Log. Program.*, 20(3):358–390, 2020.

[Greco *et al.*, 2007] Sergio Greco, Irina Trubitsyna, and Ester Zumpano. On the semantics of logic programs with preferences. *J. Artif. Intell. Res.*, 30:501–523, 2007.

[Heyninck and Strasser, 2021] Jesse Heyninck and Christian Strasser. A comparative study of assumption-based argumentative approaches to reasoning with priorities. *FLAP*, 8(3):737–808, 2021.

[Kaci *et al.*, 2018] Souhila Kaci, Leendert W. N. van der Torre, and Serena Villata. Preference in abstract argumentation. In *Proc. COMMA*, pages 405–412, 2018.

[Mailly and Rossit, 2020] Jean-Guy Mailly and Julien Rossit. Argument, I choose you! preferences and ranking semantics in abstract argumentation. In *Proc. of KR*, pages 647–651, 2020.

[Modgil and Prakken, 2013] Sanjay Modgil and Henry Prakken. A general account of argumentation with preferences. *Artif. Intell.*, 195:361–397, 2013.

[Modgil, 2009] Sanjay Modgil. Reasoning about preferences in argumentation frameworks. *Artif. Intell.*, 173(9-10):901–934, 2009.

[Papadimitriou, 1994] C. H. Papadimitriou. *Computational complexity*. Addison-Wesley, 1994.

[Prakken and Sartor, 1997] Henry Prakken and Giovanni Sartor. Argument-based extended logic programming with defeasible priorities. *J. Appl. Non Class. Logics*, 7(1):25–75, 1997.

[Sakama and Inoue, 2000] Chiaki Sakama and Katsumi Inoue. Prioritized logic programming and its application to commonsense reasoning. *Artif. Intell.*, 123(1-2), 2000.

[Silva *et al.*, 2020] Rafael Silva, Samy Sá, and João F. L. Alcântara. Semantics hierarchy in preference-based argumentation frameworks. In *COMMA*, pages 339–346, 2020.

[Toni, 2014] Francesca Toni. A tutorial on assumption-based argumentation. *Argument Comput.*, 5(1):89–117, 2014.

[Wakaki, 2015] Toshiko Wakaki. Preference-based argumentation built from prioritized logic programming. *J. Log. Comput.*, 25(2):251–301, 2015.