Abstract Argumentation Framework with Conditional Preferences

Gianvincenzo Alfano, Sergio Greco, Francesco Parisi, Irina Trubitsyna

Department of Informatics, Modeling, Electronics and System Engineering,
University of Calabria, Italy
{g.alfano, greco, fparisi, i.trubitsyna}@dimes.unical.it

Abstract

Dung's abstract Argumentation Framework (AF) has emerged as a central formalism in the area of knowledge representation and reasoning. Preferences in AF allow to represent the comparative strength of arguments in a simple yet expressive way. Preference-based AF (PAF) has been proposed to extend AF with preferences of the form $a > b$, whose intuitive meaning is that argument a is better than b. In this paper we generalize PAF by introducing *conditional preferences* of the form $a > b \leftarrow body$ that informally state that a is better than b whenever the condition expressed by *body* is true. The resulting framework, namely *Conditional Preference-based AF* (CPAF), extends the PAF semantics under three well-known preference criteria, i.e. democratic, elitist, and KTV. After introducing CPAF, we study the complexity of the verification problem (deciding whether a set of arguments is a "best" extension) as well as of the credulous and skeptical acceptance problems (deciding whether a given argument belongs to any or all "best" extensions, respectively) under multiple-status semantics (that is, complete, preferred, stable, and semi-stable semantics) for the above-mentioned preference criteria.

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 $\Omega \subseteq \mathrm{A} \times \mathrm{A}$ that specifies conflicts over arguments (if argument a attacks argument b, then b is acceptable only if a is not). Thus, an AF can be viewed 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 $\langle \mathrm{A}, \Omega \rangle = \langle \{\mathtt{a}, \mathtt{b}\}, \{(\mathtt{a}, \mathtt{b}), (\mathtt{b}, \mathtt{a})\} \rangle$ having two arguments, $\mathtt{a}$ and $\mathtt{b}$, attacking each other, there are two stable extensions, $\{\mathtt{a}\}$ and $\{\mathtt{b}\}$, and neither argument $\mathtt{a}$ nor $\mathtt{b}$ is skeptically accepted (as it will be clear after providing formal definitions, an argument is said to be skeptically accepted if it occurs in all extensions under a given semantics). To cope with such situations, a possible solution is to provide means for preferring one argument to another, as shown in the following example.

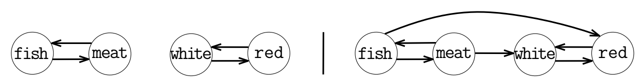

Figure 1: AF Λ_1 at the basis of the CPAF of Example 1 (left) and its rewriting Λ_1' (right).

Example 1. Consider the AF Λ_1 shown in Figure 1 (left), describing what a customer is going to have for lunch. (S)he will have either `fish` or `meat`, and will drink either `white` wine or `red` wine. Assume now that the customer expresses some preferences about the menus: if (s)he will have `meat` then would prefer to have `red` wine, whereas if (s)he will have `fish` then would prefer to have `white` wine. Intuitively, these preferences can be expressed by means of the following *conditional preferences* (CPs):

$$\mathtt{red} > \mathtt{white} \leftarrow \mathtt{meat}$$
$$\mathtt{white} > \mathtt{red} \leftarrow \mathtt{fish}$$

Λ_1 has four stable extensions (which are also preferred and semi-stable): $E_1 = \{\mathtt{fish}, \mathtt{white}\}$, $E_2 = \{\mathtt{fish}, \mathtt{red}\}$, $E_3 = \{\mathtt{meat}, \mathtt{white}\}$ and $E_4 = \{\mathtt{meat}, \mathtt{red}\}$, representing four alternative menus. However, only E_1 and E_4 are "best" extensions according to CPs expressed by the customer. $\square$

An AF with a set of conditional preferences will be called *Conditional Preference-based AF* (CPAF). The CPAF of Example 1, consisting of AF Λ_1 and the two above-mentioned conditional preferences, could be rewritten into an equivalent AF Λ_1' obtained from Λ_1 by adding the attacks $(\mathtt{meat}, \mathtt{white})$ and $(\mathtt{fish}, \mathtt{red})$ (see Figure 1 (right)). Then, AF Λ_1' has only two stable (as well as preferred and semi-stable) extensions, namely E_1 and E_4, that correspond to the best ones of the CPAF. In fact, E_2 and E_3 are no more extensions for Λ_1' as they now contain conflicting arguments.

However, in general, AF and preferences represent different pieces of knowledge, such as objective evidences and subjective beliefs, which should be clearly distinguishable. In fact, an AF represents a set of arguments and conflicts among them that leads to a set of consistent sets of argu-

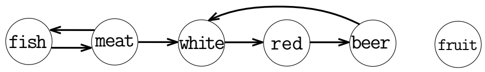

Figure 2: AF Λ_2 at the basis of the CPAF of Example 2.

ments that can be collectively accepted (i.e. the set of extensions under a given argumentation semantics) as, for instance, the alternative menus of a restaurant. In contrast, a set of preferences delivers the best extensions, e.g. best menus according to the customer's preferences as in our example.

As explained in what follows, modeling a set of conditional preferences by adding new attacks to an AF as done in Example 1 is not feasible in general.

Example 2. Consider the AF $\Lambda_2 = \langle A_2, \Omega_2 \rangle$ shown in Figure 2. Λ_2 has two preferred extensions: $E_1 = \{\texttt{fish}, \texttt{fruit}\}$ and $E_2 = \{\texttt{meat}, \texttt{red}, \texttt{fruit}\}$. Only E_2 is a stable (and semi-stable) extension. Assume that a customer expresses the following set Γ_2 of conditional preferences:

$$\texttt{fish} > \texttt{meat} \leftarrow \texttt{fruit}$$
$$\texttt{fish} > \texttt{red} \leftarrow \texttt{fruit}$$

stating that, between two menus containing `fruit`, (s)he prefers the one containing `fish` w.r.t. that containing `meat` or `red` wine. Therefore, the best extensions of the CPAF consisting of AF Λ_2 and the two conditional preferences in Γ_2 are as follows. Under the preferred semantics, the set of best preferred extensions consists of E_1 only; here E_2 is discarded in favor of E_1 that better satisfies the customer preferences. However, under stable (and semi-stable) semantics, AF Λ_2 prescribes only one extension, i.e. E_2, which now represents the best option according to the customer's preferences as it is the only one available. □

Note that the semantics of the CPAF of Example 2 cannot be represented by an equivalent AF (without preferences) as we have a situation where the best stable extensions are not contained in the best preferred extensions—this contradicts a well-known result for AF stating that every stable extension is a preferred extension (Dung 1995). That is, modifying the AF underlying a CPAF to capture preferences is not feasible in general. This is also backed by our complexity analysis entailing that CPAF cannot be reduced to AF.

AF has been extended to Preference-based Argumentation Framework (PAF) where (unconditional) preferences stating that an argument is better than another are considered. Two main approaches have been proposed in the literature in order to define PAF semantics. The first approach defines the PAF semantics in terms of that of an *auxiliary AF* (Amgoud and Cayrol 2002; Amgoud and Vesic 2014; Kaci et al. 2021). However, there are cases where this semantics may give counterintuitive results as shown next.

Example 3. Consider a PAF consisting of the AF $\Lambda_3 = \langle \{\texttt{white}, \texttt{red}, \texttt{beer}\}, \{(\texttt{white}, \texttt{red}), (\texttt{red}, \texttt{beer}), (\texttt{beer}, \texttt{white})\} \rangle$ and the (unconditional) preference $\texttt{white} > \texttt{beer}$. According to the first approach for defining PAF semantics, for the auxiliary AF $\overline{\Lambda}_3$, obtained from Λ_3 by removing attack $(\texttt{beer}, \texttt{white})$ which is conflicting with preference $\texttt{white} > \texttt{beer}$, there is only the complete extension $\{\texttt{white}, \texttt{beer}\}$. However, it is not an extension of the underlying AF Λ_3 as it is not conflict-free w.r.t. Λ_3. □

Once again, the problem is that preferences and attacks, in our opinion, describe different pieces of knowledge and should be considered separately. This is carried out by the second approach comparing extensions w.r.t. preferences defined over arguments (Amgoud and Cayrol 2002; Amgoud and Vesic 2014; Kaci et al. 2021). We follow this approach and introduce a CPAF semantics prescribing as best σ-extensions (with $\sigma \in \{\texttt{gr}, \texttt{co}, \texttt{pr}, \texttt{st}, \texttt{ss}\}$) a subset of the σ-extensions of the underlying AF that better satisfy the conditional preferences.

Contributions. Our main contributions are as follows.

- We introduce the Conditional Preference-based AF, an extension of Preference-based AF, where the underlying AF is augmented with a set of CPs. Hence, a CPAF is a triple $\langle A, \Omega, \Gamma \rangle$, where $\langle A, \Omega \rangle$ is an AF and Γ is a set of CPs.
- We propose two interpretations of conditional preferences. The *flat* interpretation only considers the preferences in Γ as they are, whereas the *closed* interpretation considers the preferences transitively obtained from Γ. We show that CPAF under closed interpretation generalizes PAF.
- We explore the complexity of the verification and credulous/skeptical acceptance problems for CPAF. The complexity of the verification problem does not depend on the flat or closed interpretation. Moreover, the complexity bounds for all the three problems for CPAF coincide with those known for PAF, though more general preferences can be expressed in CPAF.

Preliminaries

Before reviewing the Dung's framework and its generalization with preferences (PAF), we briefly recall the main complexity classes that we use (see e.g. (Papadimitriou 1994)):

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

Thus, NP^C denotes the class of problems that can be solved in polynomial time using an oracle in the class C by a non-deterministic Turing machine. It holds that $\Sigma_h^p \subseteq \Sigma_{h+1}^p \subseteq PSPACE$ and $\Pi_h^p \subseteq \Pi_{h+1}^p \subseteq PSPACE$.

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*). Different argumentation semantics have been proposed 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 and acceptable arguments w.r.t. E are defined as follows (where Λ is understood):

- $\mathit{Def}(E) = \{a \in \mathrm{A} \mid \exists b \in E \,.\, (b, a) \in \Omega\}$;

• $Acc(E) = \{a \in \mathrm{A} \mid \forall b \in \mathrm{A} \,.\, (b, a) \in \Omega \;\Rightarrow\; b \in Def(E)\}$.

To simplify the notation, we use E^+ to denote $Def(E)$.

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 $E \subseteq Acc(E)$.

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

- *complete* (co) iff it is conflict-free and $E = Acc(E)$;
- *preferred* (pr) iff it is a $\subseteq$-maximal complete extension;
- *semi-stable* (ss) iff it is a complete extension with a maximal set of decided arguments, i.e. $E \cup E^+$ is $\subseteq$-maximal;
- *stable* (st) iff it is a total complete extension ($E \cup E^+ = \mathrm{A}$);
- *grounded* (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. $\mathtt{pr}(\Lambda)$, $\mathtt{st}(\Lambda)$, $\mathtt{ss}(\Lambda)$, $\mathtt{gr}(\Lambda)$). 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. With a little abuse of notation, in the following we also use $\mathtt{gr}(\Lambda)$ to denote the grounded extension. For any AF Λ, $\mathtt{st}(\Lambda) \subseteq \mathtt{ss}(\Lambda) \subseteq \mathtt{pr}(\Lambda) \subseteq \mathtt{co}(\Lambda)$ and $\mathtt{gr}(\Lambda) \in \mathtt{co}(\Lambda)$.

Example 4. Let $\Lambda_4 = \langle \mathrm{A}_4, \Omega_4 \rangle$ be an AF where $\mathrm{A}_4 = \{\mathtt{a}, \mathtt{b}, \mathtt{c}, \mathtt{d}\}$ and $\Omega_4 = \{(\mathtt{a}, \mathtt{b}), (\mathtt{b}, \mathtt{a}), (\mathtt{a}, \mathtt{c}), (\mathtt{b}, \mathtt{c}), (\mathtt{c}, \mathtt{d}), (\mathtt{d}, \mathtt{c})\}$. The complete extensions are $E_0 = \emptyset$, $E_1 = \{\mathtt{d}\}$, $E_2 = \{\mathtt{a}, \mathtt{d}\}$ and $E_3 = \{\mathtt{b}, \mathtt{d}\}$. E_0 is the grounded extension, whereas the preferred extensions are E_2 and E_3, which are also stable and semi-stable extensions. □

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

Preference-based AF

Several works generalizing Dung's framework to handle preferences over arguments have been proposed (Amgoud and Cayrol 1998, 2002; Amgoud and Vesic 2011, 2014; Cyras 2016; Silva, Sá, and Alcântara 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.*[1]

For arguments a and b, $a > b$ means that a is better than b. Observe that also pairs in the transitive closure of $>$ are used to compare two arguments in PAF, that is, if $a > b$ and $b > c$ hold, then $a > c$ holds as well.

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

Extensions selection semantics for PAF (Amgoud and Vesic 2014; Kaci et al. 2021) handle preferences as follows. Given a PAF $\langle A, \Omega, > \rangle$, classical argumentation semantics are used to obtain the 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 determine the best extensions, corresponding to different criteria to define $\succeq$ as explained in the following definition.

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

- democratic *criterion (Amgoud and Vesic 2014): if $\forall b \in F \setminus E \; \exists a \in E \setminus F$ such that $a > b$;*
- elitist *criterion (Amgoud and Vesic 2014): if $\forall a \in E \setminus F \; \exists b \in F \setminus E$ such that $a > b$;*
- KTV *criterion (Kaci et al. 2021): 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 $F \not\succeq E$.

We use α to denote one of the three criteria in Definition 2, i.e. democratic, elitist, KTV. We often use d, e, k as shorthand for democratic, elitist, KTV, and write $\alpha \in \{d, e, k\}$.

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 $\alpha \in \{d, e, k\}$, the best σ-extensions of Δ under criterion α (denoted as $\sigma_\alpha(\Delta)$) are the extensions $E \in \sigma(\langle A, \Omega \rangle)$ such that there is no $F \in \sigma(\langle A, \Omega \rangle)$ with $F \succ E$ (under criterion α).*

Example 5. Assume to have the arguments fish, meat, white and red and consider the following three PAFs, where arguments are denoted by their initials:

$\Delta_5^1 = \langle \{\mathtt{f}, \mathtt{m}, \mathtt{r}\}, \{(\mathtt{f}, \mathtt{m}), (\mathtt{m}, \mathtt{f}), (\mathtt{f}, \mathtt{r})\}, \{\mathtt{f} > \mathtt{m}\} \rangle$,
$\Delta_5^2 = \langle \{\mathtt{f}, \mathtt{m}, \mathtt{w}\}, \{(\mathtt{f}, \mathtt{m}), (\mathtt{m}, \mathtt{f}), (\mathtt{m}, \mathtt{w})\}, \{\mathtt{f} > \mathtt{m}\} \rangle$,
$\Delta_5^3 = \langle \{\mathtt{f}, \mathtt{m}, \mathtt{r}, \mathtt{w}\}, \{(\mathtt{f}, \mathtt{m}), (\mathtt{m}, \mathtt{f}), (\mathtt{f}, \mathtt{r}), (\mathtt{m}, \mathtt{w})\}, \{\mathtt{f} > \mathtt{m}\} \rangle$.

The preferred extensions of the underlying AFs Λ_5^i ($i \in [1..3]$) obtained from Δ_5^i by ignoring the preferences are:
- $\mathtt{pr}(\Lambda_5^1) = \{E_1 = \{\mathtt{f}\},\quad E_2 = \{\mathtt{m}, \mathtt{r}\}\}$;
- $\mathtt{pr}(\Lambda_5^2) = \{E_3 = \{\mathtt{f}, \mathtt{w}\},\; E_4 = \{\mathtt{m}\}\}$;
- $\mathtt{pr}(\Lambda_5^3) = \{E_3 = \{\mathtt{f}, \mathtt{w}\},\; E_2 = \{\mathtt{m}, \mathtt{r}\}\}$.

Then, the best preferred extensions are as follows:
- $\mathtt{pr}_e(\Delta_5^1) = \mathtt{pr}_k(\Delta_5^1) = \{E_1\}$, $\mathtt{pr}_d(\Delta_1) = \{E_1, E_2\}$;
- $\mathtt{pr}_d(\Delta_5^2) = \mathtt{pr}_k(\Delta_5^2) = \{E_3\}$, $\mathtt{pr}_e(\Delta_5^2) = \{E_3, E_4\}$;
- $\mathtt{pr}_k(\Delta_5^3) = \{E_3\}$, $\mathtt{pr}_e(\Delta_5^3) = \mathtt{pr}_d(\Delta_5^3) = \{E_3, E_2\}$. □

The main difference among the above-mentioned preference criteria is that they impose different conditions to state that an extension E is preferable w.r.t. an extension F. In particular, in some situations, the democratic and elitist criteria might be too restrictive in deriving preferences among extensions. Indeed, to establish that E is preferable to F, under the democratic criterion all elements of F must be 'dominated' by some element in E, whereas under the elitist criterion all elements in E must 'dominate' at least one element in F. On the other side, the KTV criterion is less restrictive. Consider for instance the AF shown in Figure 1 (right), having the two preferred extensions $E_1 = \{\texttt{fish}, \texttt{white}\}$ and $E_4 = \{\texttt{meat}, \texttt{red}\}$, and the preference $\texttt{fish} > \texttt{meat}$. The intuitive meaning that menu E_1 should be preferable to menu E_4 is captured by the KTV criterion only. The democratic and elitist criteria state that both E_1 and E_4 are the best extensions (thus no choice is made between the two menus).

An alternative semantics 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 preorder (i.e. a reflexive and transitive relation) and $a > b$ if $a \geq b$ and $b \not\geq a$. Moreover, $E \succeq F$ 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 $c > a$, and relation $\succeq$ is reflexive ($E \succeq E$) and transitive ($E \succeq F$ and $F \succeq G$ implies $E \succeq G$). In this paper we deal with CPAF where the preference relation $\succeq$ between extensions is not transitive (as for the case of PAF where e.g. $\succeq$ is not transitive under KTV criterion), leaving the investigation of transitiveness of preferences over extensions for future work.

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, $\texttt{gr}_\alpha(\langle A, \Omega, > \rangle) = \texttt{gr}(\langle A, \Omega \rangle) \, \forall \alpha \in \{d, e, k\}$. Thus we do not consider the grounded semantics in our complexity analysis. The complexity of the verification problem as well as of the credulous and skeptical acceptance problems for PAF has been recently investigated in (Alfano et al. 2022b). 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 2018) for a survey on the complexity of AF).

AF with Conditional Preferences

In this section we extend AF with conditional preferences that allow us to express several kinds of desiderata among extensions. We first present the syntax and then give the semantics of the novel framework called Conditional Preference-based AF (CPAF).

Syntax

We augment an AF by a set of *conditional preferences* (also called preference rules) whose intuitive meaning is that an argument is better than another whenever a condition expressed by a conjunction of argument literals is satisfied. Here, an argument *literal* (or simply a literal) is either an argument a or a negated argument $\neg a$.

Definition 4. *Given an AF $\langle A, \Omega \rangle$, a conditional preference (CP) is an expression of the form*

$$a_1 > a_2 \leftarrow b_1 \wedge \cdots \wedge b_m \wedge \neg c_1 \wedge \cdots \wedge \neg c_n \qquad (1)$$

where $a_1, a_2, b_1, ..., b_m, c_1, ..., c_n$ are distinct arguments in A and $n, m \geq 0$.

For any conditional preference of the form (1), $a_1 > a_2$ is said to be the *head* of the rule, whereas the conjunction of literals $b_1 \wedge \cdots \wedge b_m \wedge \neg c_1 \wedge \cdots \wedge \neg c_n$ is the *body*. With a little abuse of notation, we often assume that the body of a rule is a set of literals (instead of a conjunction).

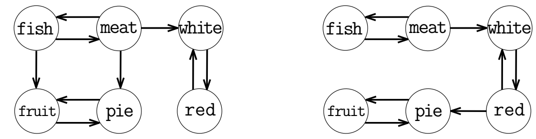

Figure 3: AFs of Example 7 (left) and Example 9 (right).

We now introduce well-formed conditional preferences.

Definition 5. *For AF $\langle A, \Omega \rangle$, a set Γ of CPs is said to be* well-formed *if there exists a function $\varphi : A \rightarrow \mathbb{N}$ such that for each CP $a > b \leftarrow body$, it holds that (i) $\varphi(a) = \varphi(b)$ and (ii) $\varphi(a) \neq \varphi(c)$ for each c (or $\neg$c) occurring in body.*

Example 6. Consider a CPAF $\Delta_6 = \langle A_6 = A_1, \Omega_6 = \Omega_1, \Gamma_6 \rangle$ obtained from the AF $\Lambda_1 = \langle A_1, \Omega_1 \rangle$ of Example 1 and the set Γ_6 of the following four CPs:

$$\begin{array}{l} \texttt{red} > \texttt{white} \leftarrow \texttt{meat} \\ \texttt{fish} > \texttt{meat} \leftarrow \texttt{white} \\ \texttt{white} > \texttt{red} \leftarrow \texttt{fish} \\ \texttt{meat} > \texttt{fish} \leftarrow \texttt{red} \end{array}$$

A possible instantiation of function φ could assign 0 to red and white, and 1 to meat and fish (or vice versa). □

The main reason for imposing well-formedness is to avoid preferences that can give counterintuitive results. For instance, consider a CPAF where the underlying AF has extensions $\{\texttt{a}, \texttt{b}\}$ and $\{\texttt{a}, \texttt{c}\}$ and the (not well-formed) preferences $\texttt{c} > \texttt{b} \leftarrow \texttt{b}$ and $\texttt{c} > \texttt{b} \leftarrow \texttt{c}$. In this situation, one would expect that $\{\texttt{a}, \texttt{c}\}$ is preferred to $\{\texttt{a}, \texttt{b}\}$. However, as it will be clear after introducing the semantics of CPAF in the next subsection, both extensions are best-extensions (under any preference criterion). On the other hand, as it will be clear in the following, using the well-formed preference $\texttt{c} > \texttt{b} \leftarrow$ we obtain the expected solution.

Throughout the paper, unless stated otherwise, we assume that conditional preferences are well-formed. Nevertheless, all our results still hold for not well-formed CPAF.

Definition 6. *A Conditional Preference-based AF (CPAF) is a triple $\langle A, \Omega, \Gamma \rangle$, where $\langle A, \Omega \rangle$ is an AF and Γ is a set of (well-formed) conditional preferences.*

Example 7. Consider AF $\Lambda_7 = \langle \mathrm{A}_7, \Omega_7 \rangle$ shown in Figure 3 (left). Let $\Delta_7 = \langle \mathrm{A}_7, \Omega_7, \Gamma_7 \rangle$, where Γ_7 consists of CPs:

$$\begin{array}{l} \texttt{red} > \texttt{white} \leftarrow \texttt{meat} \\ \texttt{white} > \texttt{red} \leftarrow \texttt{fish}. \end{array}$$

Λ_7 has three preferred extensions: $E_1 = \{\texttt{fish}, \texttt{white}, \texttt{pie}\}$, $E_2 = \{\texttt{fish}, \texttt{red}, \texttt{pie}\}$ and $E_3 = \{\texttt{meat}, \texttt{red}, \texttt{fruit}\}$ which represent possible menus. However, intuitively, we expect that the best preferred extensions according to the conditional preferences in Γ_7 are E_1 and E_3. □

Observe that the set of preference relations in the partial order $>$ defined for PAF (cf. Definition 1) can be viewed as a set of preference rules with empty body, i.e. $\{\gamma \leftarrow \mid \gamma \in >\}$.

Semantics

The meaning of a CPAF $\langle A, \Omega, \Gamma\rangle$ w.r.t. a given argumentation semantics $\sigma \in \{\texttt{gr}, \texttt{co}, \texttt{pr}, \texttt{st}, \texttt{ss}\}$ is given by considering the extensions that better satisfy Γ among the σ-extensions of the underlying AF $\langle A, \Omega\rangle$. This is carried out by extending the PAF comparison criteria between extensions (i.e. democratic, elitist and KTV of Definition 2) according to two different interpretations of the preference rules, that are *flat* and *closed* interpretations. As discussed in what follows, differently from the flat interpretation, the closed interpretation deals with the closure of Γ.

Before providing the semantics of CPAF under flat and closed interpretation, we introduce some notation. For any conditional preference $\gamma = a_1 > a_2 \leftarrow b_1 \wedge \cdots \wedge b_m \wedge \neg c_1 \wedge \cdots \wedge \neg c_n$ and (conflict-free) set of argument E, we say that E *satisfies the body of* γ (and write $E \models b_1 \wedge \cdots \wedge b_m \wedge \neg c_1 \wedge \cdots \wedge \neg c_n$) iff $\{b_1, ..., b_m\} \subseteq E$ and $\{c_1, ..., c_n\} \subseteq E^+$, that is the arguments that positively (resp. negatively) occur in the body of γ belong to E (resp. E^+).

Flat interpretation. The next definition introduces the democratic, elitist and KTV preference criteria for CPAF.

Definition 7. *Given a CPAF $\langle A, \Omega, \Gamma\rangle$, for $E, F \subseteq A$ with $E \neq F$, we have that $E \succeq F$ under*

- democratic (d) *criterion:*
  if $\forall\, b \in F \setminus E\ \exists\, a \in E \setminus F$ and $\exists\, a > b \leftarrow body \in \Gamma$ such that $E \models body$ and $F \models body$;
- elitist (e) *criterion:*
  if $\forall\, a \in E \setminus F\ \exists\, b \in F \setminus E$ and $\exists\, a > b \leftarrow body \in \Gamma$ such that $E \models body$ and $F \models body$;
- KTV (k) *criterion:*
  if $\forall\, a, b \in A\ \nexists\, a > b \leftarrow body \in \Gamma$ such that $a \in F \setminus E$, $b \in E \setminus F$, $E \models body$, and $F \models body$.

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

Best extensions under flat interpretation are as follows.

Definition 8. *Given a CPAF $\Delta = \langle A, \Omega, \Gamma\rangle$, a semantics $\sigma \in \{\texttt{gr}, \texttt{co}, \texttt{pr}, \texttt{st}, \texttt{ss}\}$, and a criterion $\alpha \in \{d, e, k\}$, the best σ-extensions of Δ under criterion α (denoted as $\sigma_\alpha(\Delta)$) are the extensions $E \in \sigma(\langle A, \Omega\rangle)$ such that there is no $F \in \sigma(\langle A, \Omega\rangle)$ with $F \succ E$ (under criterion α).*

Example 8. Continuing with Example 7, we have that $E_1 \succ E_2$, $E_1 \not\succeq E_3$, $E_2 \not\succeq E_3$, $E_3 \not\succeq E_1$ and $E_3 \not\succeq E_2$ under democratic, elitist and KVT criteria. Thus, E_1 and E_3 are the best preferred extensions of the CPAF Δ_7 of Example 7, that is, $\texttt{pr}_\alpha(\Delta_7) = \{E_1, E_3\}$ with $\alpha \in \{d, e, k\}$. $\square$

Example 9. Consider the AF $\Lambda_9 = \langle \mathrm{A}_9, \Omega_9\rangle$ shown in Figure 3 (right). Let Γ_9 be the set of the following CPs:

$$\texttt{fish} > \texttt{meat} \leftarrow \texttt{fruit}$$
$$\texttt{white} > \texttt{red} \leftarrow \texttt{fish}$$

For the CPAF $\Delta_9 = \langle \mathrm{A}_9, \Omega_9, \Gamma_9\rangle$, there are four preferred (and stable/semi-stable) extensions: $E_1 = \{\texttt{fish}, \texttt{white}, \texttt{pie}\}$, $E_2 = \{\texttt{fish}, \texttt{white}, \texttt{fruit}\}$, $E_3 = \{\texttt{fish}, \texttt{red}, \texttt{fruit}\}$, and $E_4 = \{\texttt{meat}, \texttt{red}, \texttt{fruit}\}$. We have that $E_2 \succ E_3$ and $E_3 \succ E_4$ under democratic, elitist and KTV criteria, whereas $E_1 \succ E_3$ and $E_2 \succ E_4$ under KTV criterion. Thus, E_1 and E_2 are the best preferred (and stable/semi-stable) extensions under democratic, elitist and KTV criteria. $\square$

Finally, considering the examples in the Introduction, we have that for the CPAF Δ_1 of Example 1, $E_1 \succ E_2$ and $E_4 \succ E_3$ under democratic, elitist and KVT criteria, and thus $\sigma_\alpha(\Delta_1) = \{E_1, E_4\}$ with $\sigma \in \{\texttt{st}, \texttt{pr}, \texttt{ss}\}$ and $\alpha \in \{d, e, k\}$; the result $\sigma_\alpha(\Delta_1) = \{E_1, E_4\}$ also holds for CPAF Δ_6 of Example 6 that extends Δ_1 with two additional CPs. Moreover, for the CPAF $\Delta_2 = \langle A_2, \Omega_2, \Gamma_2\rangle$ of Example 2, since $E_1 \succ E_2$ under democratic, elitist and KVT criteria, we have that $\texttt{pr}_\alpha(\Delta_2) = \{E_1\}$ and, since $\texttt{st}(\langle A_2, \Omega_2\rangle) = \{E_2\}$, $\texttt{st}_\alpha(\Delta_2) = \texttt{ss}_\alpha(\Delta_2) = \texttt{st}(\langle A_2, \Omega_2\rangle) = \{E_2\}$ with $\alpha \in \{d, e, k\}$.

Closed interpretation. The CPAF with flat interpretation (of preference rules) does not generalize the PAF, in the sense that the semantics of a CPAF $\langle A, \Omega, \Gamma\rangle$ where Γ consists of unconditional preferences (preference rules with empty body) may be not equivalent to considering a strict partial order over arguments as in PAF. Therefore, we now propose a different semantics, called *closed interpretation*, that generalizes that of PAF.

The closed interpretation assumes that Γ denotes all dependencies logically implied by it. To this end we introduce the closure of Γ, defined iteratively as follows:

$$\begin{aligned}\Gamma^* = \Gamma \cup \{ a_1 > a_3 \leftarrow body_1 \wedge body_2 \mid \\ \exists\, a_1 > a_2 \leftarrow body_1 \in \Gamma^* \wedge \\ \exists\, a_2 > a_3 \leftarrow body_2 \in \Gamma^* \}.\end{aligned}$$

Notice that the size of Γ^* may be exponentially larger than that of Γ.

We can build on the flat interpretation to define the semantics of a CPAF under closed interpretation as follows.

Definition 9. *Let $\Delta = \langle A, \Omega, \Gamma\rangle$ be a CPAF, $\sigma \in \{\texttt{gr}, \texttt{co}, \texttt{pr}, \texttt{st}, \texttt{ss}\}$ an argumentation semantics, and $\alpha \in \{d, e, k\}$ a preference criterion.*

1. *For $E, F \subseteq A$, we say that E is preferred to F under closed interpretation and criterion α (denoted as $E \succeq^* F$) iff $E \succeq F$ in the CPAF $\langle A, \Omega, \Gamma^*\rangle$ under flat interpretation and criterion α.*
   Moreover, $E \succ^* F$ if $E \succeq^* F$ and $F \not\succeq^* E$.
2. *The best σ-extensions of Δ under criterion α and closed interpretation (denoted as $\sigma^*_\alpha(\Delta)$) are the extensions $E \in \sigma(\langle A, \Omega\rangle)$ such that there is no $F \in \sigma(\langle A, \Omega\rangle)$ with $F \succ^* E$ (under criterion α).*

As stated below, the best extensions under closed interpretation are obtained by just taking Γ^* instead of Γ.

Fact 1. *For any CPAF $\Delta = \langle A, \Omega, \Gamma\rangle$, semantics $\sigma \in \{\texttt{gr}, \texttt{co}, \texttt{pr}, \texttt{st}, \texttt{ss}\}$ and criterion $\alpha \in \{d, e, k\}$, it holds that $\sigma^*_\alpha(\langle A, \Omega, \Gamma\rangle) = \sigma_\alpha(\langle A, \Omega, \Gamma^*\rangle)$.*

The next example shows that a CPAF with unconditional preferences under flat interpretation behaves differently from the corresponding PAF, whereas the closed interpretation gives results equal to those of the PAF.

Example 10. Consider the CPAF $\Delta_{10} = \langle \mathrm{A}_{10} = \{\texttt{red}, \texttt{white}, \texttt{beer}\}, \Omega_{10} = \{(\texttt{red}, \texttt{beer}), (\texttt{beer}, \texttt{red})\}, \Gamma_{10} = \{\ \texttt{red} > \texttt{white} \leftarrow,\ \texttt{white} > \texttt{beer} \leftarrow\ \}\rangle$. For the AF $\langle A_{10}, \Omega_{10}\rangle$, there are two preferred and stable extensions, $E_1 = \{\texttt{red}, \texttt{white}\}$ and $E_2 = \{\texttt{beer}, \texttt{white}\}$. Under flat interpretation, we have that $\mathtt{pr}_\alpha(\Delta_{10}) = \mathtt{pr}(\langle A_{10}, \Omega_{10}\rangle) = \{E_1, E_2\}$ for all $\alpha \in \{d, e, k\}$, whereas under closed interpretation we have that $\mathtt{pr}^*_\alpha(\Delta_{10}) = \{E_1\}$. Considering the corresponding PAF $\Delta'_{10} = \langle \mathrm{A}_{10}, \Omega_{10}, \{\mathtt{red} > \mathtt{white},\ \mathtt{white} > \mathtt{beer}\}\rangle$, we have $\mathtt{pr}_\alpha(\Delta') = \{E_1\}$, that is, the closed interpretation is the one that directly models the PAF semantics. □

It is worth mentioning that for the CPAFs of Example 1 and Example 2 the closed and flat interpretation coincide.

As for the possible situations in which it would be better to use the flat or the closed interpretation, an important point is that it may depend on the context in which the user operates and on the user's familiarity with the use of preferences. In fact, on one hand, the closed interpretation allows for a more compact representation of preferences, though (transitive) preferences not immediately visible to the user are considered in the process. On the other hand, the flat interpretation provides the user with explicit control over the set of preferences to be considered, but transitive preferences have to be given explicitly, otherwise they will ignored.

As stated next, if Γ is well-formed then Γ^* is well-formed.

Theorem 1. *For any set Γ of well-formed conditional preferences, it holds that Γ^* is well-formed as well.*

Properties of CPAF

The following proposition states that any conditional preference having an head argument occurring in the body does not play any role (under flat or closed interpretation). Note that this kind of conditional preferences are not well-formed. That is, Definition 5 avoids using useless CPs.

Proposition 1. *Let $\Delta = \langle A, \Omega, \Gamma\rangle$ be such that $\gamma = a_1 > a_2 \leftarrow b_1, \cdots, b_m, \neg c_1, \cdots, \neg c_n$ belongs to Γ (resp. Γ^*) and $\{a_1, a_2\} \cap \{b_1, \ldots, b_m, c_1, \ldots, c_n\} \neq \emptyset$. Then, for any semantics $\sigma \in \{\mathtt{gr}, \mathtt{co}, \mathtt{pr}, \mathtt{st}, \mathtt{ss}\}$ and criterion $\alpha \in \{d, e, k\}$, it holds that $\sigma_\alpha(\Delta) = \sigma_\alpha(\langle A, \Omega, \Gamma \setminus \{\gamma\}\rangle$ (resp. $\sigma^*_\alpha(\Delta) = \sigma_\alpha(\langle A, \Omega, \Gamma^* \setminus \{\gamma\}\rangle$.*

The next proposition states that under closed interpretation CPAF semantics extend PAF semantics, and this holds under flat interpretation if unconditional preferences representing the closure of the PAF preferences are considered.

Proposition 2. *Let $\Delta = \langle A, \Omega, >\rangle$ be a PAF, $\sigma \in \{\mathtt{gr}, \mathtt{co}, \mathtt{pr}, \mathtt{st}, \mathtt{ss}\}$ a semantics, and $\alpha \in \{d, e, k\}$ a criterion. It holds that $E \in \sigma_\alpha(\Delta)$ iff $E \in \sigma^*_\alpha(\langle A, \Omega, \Gamma\rangle)$, where $\Gamma = \{\gamma \leftarrow \mid \gamma \in >\}$.*

As observed earlier, for any AF Λ, $\mathtt{st}(\Lambda) \subseteq \mathtt{ss}(\Lambda) \subseteq \mathtt{pr}(\Lambda)$. However, as stated next, in general, the set of the best stable (resp. semi-stable) extensions of a CPAF is not a proper subset of the set of the best preferred extensions; the result holds irrespective of the interpretation (flat/closed) and preference criterion (democratic/elitist/KTV).

Proposition 3. Let $\alpha \in \{d, e, k\}$ be a criterion, then:

- There exists a CPAF Δ such that $\mathtt{st}_\alpha(\Delta) \not\subseteq \mathtt{pr}_\alpha(\Delta)$ and $\mathtt{st}^*_\alpha(\Delta) \not\subseteq \mathtt{pr}^*_\alpha(\Delta)$.
- There exists a CPAF Δ such that $\mathtt{ss}_\alpha(\Delta) \not\subseteq \mathtt{pr}_\alpha(\Delta)$ and $\mathtt{ss}^*_\alpha(\Delta) \not\subseteq \mathtt{pr}^*_\alpha(\Delta)$.
- For every CPAF Δ, it holds that:
  $\mathtt{st}_\alpha(\Delta) \neq \emptyset \Rightarrow \mathtt{st}_\alpha(\Delta) = \mathtt{ss}_\alpha(\Delta)$ and
  $\mathtt{st}^*_\alpha(\Delta) \neq \emptyset \Rightarrow \mathtt{st}^*_\alpha(\Delta) = \mathtt{ss}^*_\alpha(\Delta)$.

Observe that the result of the last item in Proposition 3 entails that, for every CPAF Δ, $\mathtt{st}_\alpha(\Delta) \subseteq \mathtt{ss}_\alpha(\Delta)$ and $\mathtt{st}^*_\alpha(\Delta) \subseteq \mathtt{ss}^*_\alpha(\Delta)$, analogously to what holds for AF.

Proposition 3 suggests that preferences cannot be represented in (classical) AF in general, as we showed that there are situations where the best stable extensions are not contained in the best preferred extensions (a situation that is not reflected in AF). This will be also confirmed by the complexity results showing that the considered problems for CPAF are generally harder than those for AF.

We conclude this section by presenting a proposition stating that, if we are interested to compare two extensions under the flat interpretation, then we can focus on a restricted set of unconditional preferences, that is, we can refer to a CPAF with a restricted set of preferences. An analogous result holds also for the closed interpretation. Before formalizing the results, we introduce some notation. For any CPAF $\langle A, \Omega, \Gamma\rangle$ and set $E \subseteq A$, we use Γ_E to denote the set $\{a_1 > a_2 \leftarrow \mid \exists\, a_1 > a_2 \leftarrow \mathit{body} \in \Gamma \wedge E \models \mathit{body}\}$ of unconditional preferences derived from Γ and E.

Proposition 4. *For any CPAF $\Delta = \langle A, \Omega, \Gamma\rangle$, semantics $\sigma \in \{\mathtt{gr}, \mathtt{co}, \mathtt{pr}, \mathtt{st}, \mathtt{ss}\}$, criterion $\alpha \in \{d, e, k\}$, and pair $E, F \subseteq A$ with $E \neq F$, it holds that:*

- *$E \succeq F$ w.r.t. Δ iff $E \succeq F$ w.r.t. $\langle A, \Omega, \Gamma_E \cap \Gamma_F\rangle$; and*
- *$E \succeq^* F$ w.r.t. Δ iff $E \succeq F$ w.r.t. $\langle A, \Omega, (\Gamma_E \cap \Gamma_F)^*\rangle$.*

Since $\Gamma_E \cap \Gamma_F$ consists of unconditional preferences only, and the closure $(\Gamma_E \cap \Gamma_F)^*$ can be computed in polynomial time, the results of Proposition 4 entail that comparing two extensions under flat or closed interpretation is polynomial too. That is, relying on the possible exponentially large set Γ^* (as in Definition 9) is not needed. This result is formally stated in Proposition 7 in the next section.

Complexity of CPAF

We investigate the complexity of some fundamental problems in the proposed framework, that are the verification problem as well as the credulous and skeptical acceptance problems. We start by introducing some results that will be useful to characterize the complexity of these problems.

Theorem 2. *The problem of checking whether a set of conditional preferences is well-formed is polynomial time.*

The following proposition considers the satisfaction of CPs by extensions which are related by subset inclusion.

Proposition 5. *Let $\Lambda = \langle A, \Omega\rangle$ be an AF, $E, F \in \mathtt{co}(\Lambda)$ two complete extensions of Λ, and $\gamma = a_1 > a_2 \leftarrow \mathit{body}$ a CP. It holds that, if $E \subseteq F$ and $E \models \mathit{body}$, then $F \models \mathit{body}$.*

The following proposition states that, irrespective of the flat or closed interpretation, best complete and grounded semantics for CPAF 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 6. *For any CPAF $\Delta = \langle A, \Omega, \Gamma\rangle$ it holds that:*

i) $\texttt{co}_e(\Delta) = \texttt{co}^*_e(\Delta) = \texttt{gr}(\langle A, \Omega\rangle)$,

ii) $\texttt{co}_d(\Delta) = \texttt{pr}_d(\Delta)$ *and* $\texttt{co}^*_d(\Delta) = \texttt{pr}^*_d(\Delta)$*, and*

iii) $\texttt{gr}(\langle A, \Omega\rangle) \in \texttt{co}_k(\Delta)$ *and* $\texttt{gr}(\langle A, \Omega\rangle) \in \texttt{co}^*_k(\Delta)$.

The next proposition states that comparing two extensions under flat or closed interpretation is polynomial.

Proposition 7. *Let $\Delta = \langle A, \Omega, \Gamma\rangle$ be a CPAF, $\sigma \in \{\texttt{gr}, \texttt{co}, \texttt{pr}, \texttt{st}, \texttt{ss}\}$ a semantics, and $\alpha \in \{d, e, k\}$ a criterion. For any pair of extensions $E, F \in \sigma(\langle A, \Omega\rangle)$, the problem of deciding whether $E \succeq F$ (resp. $E \succeq^* F$) under criterion α and flat (resp. closed) interpretation is polynomial.*

We are now ready to present the complexity of the verification problem for CPAF.

Verification Problem

The verification problem for CPAF under flat interpretation, denoted as Ver_{σ_α} with $\sigma \in \{\texttt{gr}, \texttt{co}, \texttt{pr}, \texttt{st}, \texttt{ss}\}$ and $\alpha \in \{d, e, k\}$, extends that for AF (as well as that for PAF) by considering the best σ_α-extensions of a CPAF. That is, given a CPAF $\Delta = \langle \mathrm{A}, \Omega, \Gamma\rangle$, Ver_{σ_α} is the problem of deciding whether a set of arguments $S \subseteq \mathrm{A}$ belongs to $\sigma_\alpha(\Delta)$. The definition for the closed interpretation is the same except that σ^*_α is considered instead of σ_α, i.e. we check if $S \in \sigma^*_\alpha(\Delta)$.

Recall that for the grounded semantics the verification problem for CPAF is equivalent to that for AF, which can be checked in polynomial time (Dung 1995). The following theorem characterizes the complexity of the verification problem for the different combinations of multiple-status semantics $\sigma \in \{\texttt{co}, \texttt{pr}, \texttt{st}, \texttt{ss}\}$, criterion $\alpha \in \{d, e, k\}$, and flat or closed interpretations.[2]

Theorem 3. *Ver_{σ_α} and $Ver_{\sigma^*_\alpha}$ are:*

- *in P for $\sigma_\alpha = \texttt{co}_e$;*
- *co*NP*-complete for $\sigma_\alpha \in \{\texttt{co}_k, \texttt{st}_d, \texttt{st}_e, \texttt{st}_k, \texttt{pr}_d\}$; and*
- *Π^p_2-complete for $\sigma_\alpha \in \{\texttt{pr}_e, \texttt{pr}_k, \texttt{ss}_d, \texttt{ss}_e, \texttt{ss}_k\}$.*

Thus, the complexity of the verification problem for CPAF does not depend on the flat or closed interpretation and coincides with that of PAF where only unconditional preferences are considered. Notably, these complexity results coincide with those for PAF. The results of Theorem 3 are also useful to analyze the complexity of the credulous and skeptical acceptance problems, which is addressed next.

Credulous and Skeptical Acceptance

The credulous and skeptical acceptance problems for CPAF are defined as expected by considering the best σ_α-extensions (or best σ^*_α-extensions). More formally, given a CPAF $\Delta = \langle \mathrm{A}, \Omega, \Gamma\rangle$ and a goal argument $g \in \mathrm{A}$, the credulous (resp. skeptical) acceptance problem under flat interpretation, denoted as CA_{σ_α} (resp. SA_{σ_α}), consists in deciding whether g belongs to any (resp. every) extension in $\sigma_\alpha(\Delta)$; for the closed interpretation, we consider σ^*_α instead of σ_α, with $\sigma \in \{\texttt{gr}, \texttt{co}, \texttt{pr}, \texttt{st}, \texttt{ss}\}$ and $\alpha \in \{d, e, k\}$.

The complexity of the credulous acceptance problem under flat and closed interpretation is as follows.

[2] The result for $\texttt{co}_d$ is not stated since $\texttt{co}_d(\Delta) = \texttt{pr}_d(\Delta)$ and $\texttt{co}^*_d(\Delta) = \texttt{pr}^*_d(\Delta)$, as shown in Proposition 6.

Theorem 4. *CA_{σ_α} and $CA_{\sigma^*_\alpha}$ are:*

- *in P for $\sigma_\alpha = \texttt{co}_e$;*
- *Σ^p_2-complete for $\sigma_\alpha \in \{\texttt{co}_k, \texttt{st}_d, \texttt{st}_e, \texttt{st}_k, \texttt{pr}_d\}$;*
- *Σ^p_2-hard and in Σ^p_3 for $\sigma_\alpha \in \{\texttt{pr}_e, \texttt{pr}_k, \texttt{ss}_d, \texttt{ss}_e, \texttt{ss}_k\}$.*

Finally, we consider the skeptical acceptance problem.

Theorem 5. *SA_{σ_α} and $SA_{\sigma^*_\alpha}$ are:*

- *in P for $\sigma_\alpha \in \{\texttt{co}_e, \texttt{co}_k\}$;*
- *Π^p_2-complete for $\sigma_\alpha \in \{\texttt{st}_d, \texttt{st}_e, \texttt{st}_k, \texttt{pr}_d\}$; and*
- *Π^p_2-hard and in Π^p_3 for $\sigma_\alpha \in \{\texttt{pr}_e, \texttt{pr}_k, \texttt{ss}_d, \texttt{ss}_e, \texttt{ss}_k\}$.*

The results of this section show that the complexity of the three considered problems for CPAF generally increases of one level in the polynomial hierarchy w.r.t that of AF. Moreover, the complexity bounds obtained for CPAF problems are the same as those known for PAF (Alfano et al. 2022b), though more general preferences can be expressed in CPAF.

Related Work

Preferences have been extensively studied in AI and several formalisms have been proposed to express and reason with different kinds of preferences (see e.g. (Brafman and Domshlak 2009)). Conditional Preference networks (CP-nets) (Boutilier et al. 2004; Rossi, Venable, and Walsh 2004) are among the most studied formalisms (Allen et al. 2017; Goldsmith et al. 2008; Lukasiewicz and Malizia 2019, 2022) and allow to express sets of conditional *ceteris paribus* (i.e. all else being equal) preference statements. For instance, one could express statements of the form "I prefer red wine to white wine whenever I have meat, all else being equal in the rest of the meal". This statement asserts that, given two meals that only differ in the kind of wine and both containing meat, meal with red wine is preferable to meal with white wine. This kind of reasoning is related to CPAF semantics in the sense that our preference rules allow for comparing extensions (i.e. alternatives) that share some parts (i.e. conditions of rules' body that are satisfied by extensions).

Previous works on embedding preferences in AF has concerned the study of PAF, where preferences are of the form $a > b$ and define a strict partial order over arguments. The semantics of PAF has been defined by considering either (i) a mapping to AF (Amgoud and Cayrol 2002; Amgoud and Vesic 2014; Kaci et al. 2021) or (ii) selection of the best extensions w.r.t. the given preferences (Amgoud and Vesic 2014; Wakaki 2015; Kaci et al. 2021). Conditional preferences in AF have been recently investigated in (Bernreiter, Dvořák, and Woltran 2022) for PAFs whose semantics is defined by using the first approach. In this paper we follow the second approach as the first one could give less intuitive results (cf. Example 3). A combination of the two approaches is the *rich PAF* (Amgoud and Vesic 2014).

In Epistemic AF (EAF) (Sakama and Son 2020), a form of conditional preferences can be expressed by using *epistemic constraints*, that is constraints defined by logical formulae containing epistemic atoms of the form $\mathbf{K}(\varphi)$ and $\mathbf{M}(\varphi)$, where φ is a propositional formula built over the labelling of arguments. Intuitively, $\mathbf{K}(\varphi)$ (resp. $\mathbf{M}(\varphi)$) states that an agent believes that φ is certainly true (resp. possibly true). Thus, a preference $x > y$ can be expressed as

$\mathbf{K}(\mathbf{in}(y) \Rightarrow \mathbf{in}(x))$, meaning that argument x should be accepted whenever the other argument y of lower preference is accepted. A preference of the form $x > y \leftarrow z$ can be expressed as $\mathbf{M}(\mathbf{in}(z)) \Rightarrow \mathbf{K}(\mathbf{in}(y) \Rightarrow \mathbf{in}(x))$. However, EAF and CPAF are significantly different. Indeed, while conditional preferences in EAF express preferences over justification states, preferences in CPAF express formulae to compare extensions. Furthermore, although epistemic constraints are of great interest because they allow to express not only preferences but also general (epistemic) acceptance conditions, according to the complexity results presented in the paper, CPAF preferences cannot be fully captured in EAF.

Preferences can also be expressed in value-based AFs (VAFs) (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. Building on Dung acceptability semantics, *Extended AF* (EAF) (Modgil 2009) provides an unifying treatment of PAF and VAF. By incorporating attacks towards attacks (i.e. second-order attacks), EAFs also accommodate argumentation based reasoning with possibly contradictory preferences. In this regard, several frameworks extending AF with higher-order relations (e.g. second-order attacks) have been proposed (Cayrol, Cohen, and Lagasquie-Schiex 2021). However, they can be rewritten equivalently in terms of (meta-)AFs, leading to the same complexity of that of AF (Villata et al. 2012; Gottifredi et al. 2018). Notably, this is not the case of CPAF where the complexity generally increases w.r.t. that of AF. In (Dunne et al. 2011; Coste-Marquis et al. 2012) weights are associated with attacks, and new semantics extending the classical ones are proposed. Recently, (Mailly and Rossit 2020) has investigated to what extent preferences can be blended with ranking semantics. In this regard, it would be interesting to explore the connection between ranking semantics (Bonzon et al. 2016) and CPs.

There has been extensive research on rule-based systems with prioritized rules and in particular on Answer Set programming where preference rules are used to filter out the best models (Brewka 1989; Brewka and Eiter 1999; Delgrande, Schaub, and Tompits 2003; Eiter et al. 2003; Sakama and Inoue 2000; Schaub and Wang 2001; Brewka, Niemelä, and Truszczynski 2003; Greco, Trubitsyna, and Zumpano 2007; Brewka et al. 2015).

The work of (Prakken and Sartor 1997) is one of the first attempt that investigates the application of priorities in defeasible rule systems to first define a preference order between arguments and then using that order to remove unwanted attacks. (Prakken 2010) and (Modgil and Prakken 2013, 2014) proposed ASPIC$^+$, a rich framework for structured argumentation with prioritized rules with several systems of preference orders between arguments. Preferences in structured argumentation formalisms are typically used to resolve attacks into defeats between arguments (Cyras et al. 2018; Garcia, Prakken, and Simari 2020; Alfano et al. 2021a; Heyninck and Strasser 2021). All these proposals mostly deal with systems with unconditional preferences. Recently, (Dung, Thang, and Son 2019) has proposed a framework for dealing with conditional preferences in structured argumentation, by defining defeasible knowledge bases with conditional preferences over a rule-based system. An argument-based framework for multi-criteria decision-making based on conditional rules has been introduced in (Brarda, Tamargo, and García 2019). Finally, an approach to handle multiple argument preference criteria in argumentation-based decision support systems has been proposed in (Teze et al. 2020). All the above-mentioned works deal with structured argumentation. To the best of our knowledge, this is the first paper proposing CPs in abstract argumentation.

Conclusion and Future Work

We have introduced the CPAF framework, an extension of PAF framework where preferences between arguments are conditioned to the acceptance of other arguments. In this setting, after studying some properties of CPAF, we investigated the complexity of three fundamental problems for CPAF (verification, credulous and skeptical acceptance).

We envisage several interesting directions for future work. Besides exploring the relationships between CPAF and rich PAF (Amgoud and Vesic 2014), as well as with ranking semantics for AF (Bonzon et al. 2016; Mailly and Rossit 2020), we plan to further investigate other preference criteria to compare extensions such as the ones defined for comparing ASP models (Sakama and Inoue 2000; Brewka, Niemelä, and Truszczynski 2003).

Conditional preferences allow discarding extensions that do not meet user preferences defined over arguments. Related to this, *extension removal* (Baumann and Brewka 2019) is concerned with manipulating an AF so that some undesired extensions given in input are removed. Hence, conditional preferences could be simulated by forcing the removal of all not-best extensions. In this regard, it would be interesting to investigate in more detail the relationship between CPAF and extension removal in future work.

We also plan to investigate conditional preferences in other argumentation frameworks whose semantics is closely related to that of AF (Alfano et al. 2020c, 2023), as well as in incomplete AF (Fazzinga, Flesca, and Furfaro 2020; Alfano et al. 2022a), probabilistic AF (Fazzinga, Flesca, and Parisi 2015; Alfano et al. 2020a), and AF with constraints (Alfano et al. 2021b). Another interesting direction for future work is exploring CPAF in a dynamic setting (Alfano et al. 2018; Alfano, Greco, and Parisi 2018; Alfano and Greco 2021; Alfano et al. 2020b; Niskanen and Järvisalo 2020; Alfano, Greco, and Parisi 2021), where objective evidences (underlying AF) and subjective beliefs (conditional preferences) may change over the time.

Finally, following e.g. (Ganzer-Ripoll et al. 2019) where some relationships between social choice theory (Kelly 1977; Fishburn 1972; Gärdenfors 1979) and argumentation have been investigated, it would be interesting to study the formal connection between preference extensions in social choice theory and preference criteria of (C)PAF in order to explore new preference criteria for argumentation.

Acknowledgments

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

References

Alfano, G.; Calautti, M.; Greco, S.; Parisi, F.; and Trubitsyna, I. 2020a. Explainable Acceptance in Probabilistic Abstract Argumentation: Complexity and Approximation. In *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning*, 33–43.

Alfano, G.; Cohen, A.; Gottifredi, S.; Greco, S.; Parisi, F.; and Simari, G. 2020b. Dynamics in Abstract Argumentation Frameworks with Recursive Attack and Support Relations. In *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI)*, 577–584.

Alfano, G.; and Greco, S. 2021. Incremental Skeptical Preferred Acceptance in Dynamic Argumentation Frameworks. *IEEE Intell. Syst.*, 36(2): 6–12.

Alfano, G.; Greco, S.; and Parisi, F. 2018. A meta-argumentation approach for the efficient computation of stable and preferred extensions in dynamic bipolar argumentation frameworks. *Intelligenza Artificiale*, 12(2): 193–211.

Alfano, G.; Greco, S.; and Parisi, F. 2021. Incremental Computation in Dynamic Argumentation Frameworks. *IEEE Intell. Syst.*, 36(6): 80–86.

Alfano, G.; Greco, S.; Parisi, F.; Simari, G. I.; and Simari, G. R. 2018. An Incremental Approach to Structured Argumentation over Dynamic Knowledge Bases. In *Proceedings of International Conference on Principles of Knowledge Representation and Reasoning (KR)*, 78–87.

Alfano, G.; Greco, S.; Parisi, F.; Simari, G. I.; and Simari, G. R. 2021a. Incremental computation for structured argumentation over dynamic DeLP knowledge bases. *Artif. Intell.*, 300: 103553.

Alfano, G.; Greco, S.; Parisi, F.; and Trubitsyna, I. 2020c. On the Semantics of Abstract Argumentation Frameworks: A Logic Programming Approach. *Theory Pract. Log. Program.*, 20(5): 703–718.

Alfano, G.; Greco, S.; Parisi, F.; and Trubitsyna, I. 2021b. Argumentation Frameworks with Strong and Weak Constraints: Semantics and Complexity. In *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence*, 6175–6184.

Alfano, G.; Greco, S.; Parisi, F.; and Trubitsyna, I. 2022a. Incomplete Argumentation Frameworks: Properties and Complexity. In *Proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence*, 5451–5460.

Alfano, G.; Greco, S.; Parisi, F.; and Trubitsyna, I. 2022b. On Preferences and Priority Rules in Abstract Argumentation. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence*, 2517–2524.

Alfano, G.; Greco, S.; Parisi, F.; and Trubitsyna, I. 2023. On acceptance conditions in abstract argumentation frameworks. *Information Sciences*, 625: 757–779.

Allen, T. E.; Goldsmith, J.; Justice, H. E.; Mattei, N.; and Raines, K. 2017. Uniform Random Generation and Dominance Testing for CP-Nets. *J. Artif. Intell. Res.*, 59: 771–813.

Amgoud, L.; and Cayrol, C. 1998. On the Acceptability of Arguments in Preference-based Argumentation. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, 1–7.

Amgoud, L.; and Cayrol, C. 2002. Inferring from Inconsistency in Preference-Based Argumentation Frameworks. *J. Autom. Reason.*, 29(2): 125–169.

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

Amgoud, L.; and Vesic, S. 2014. Rich preference-based argumentation frameworks. *Int. J. Approx. Reason.*, 55(2): 585–606.

Baumann, R.; and Brewka, G. 2019. Extension Removal in Abstract Argumentation - An Axiomatic Approach. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence*, 2670–2677.

Bench-Capon, T. J. M. 2003. Persuasion in Practical Argument Using Value-based Argumentation Frameworks. *J. Log. Comput.*, 13(3): 429–448.

Bernreiter, M.; Dvorák, W.; and Woltran, S. 2022. Abstract Argumentation with Conditional Preferences. In *Proc. of COMMA*, 92–103.

Bonzon, E.; Delobelle, J.; Konieczny, S.; and Maudet, N. 2016. A Comparative Study of Ranking-Based Semantics for Abstract Argumentation. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, 914–920.

Boutilier, C.; Brafman, R. I.; Domshlak, C.; Hoos, H. H.; and Poole, D. 2004. CP-nets: A Tool for Representing and Reasoning with Conditional Ceteris Paribus Preference Statements. *J. Artif. Intell. Res.*, 21: 135–191.

Brafman, R. I.; and Domshlak, C. 2009. Preference Handling - An Introductory Tutorial. *AI Mag.*, 30(1): 58–86.

Brarda, M. E. B.; Tamargo, L. H.; and García, A. J. 2019. An approach to enhance argument-based multi-criteria decision systems with conditional preferences and explainable answers. *Expert Syst. Appl.*, 126: 171–186.

Brewka, G. 1989. Preferred Subtheories: An Extended Logical Framework for Default Reasoning. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, 1043–1048.

Brewka, G.; Delgrande, J. P.; Romero, J.; and Schaub, T. 2015. asprin: Customizing Answer Set Preferences without a Headache. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, 1467–1474.

Brewka, G.; and Eiter, T. 1999. Preferred Answer Sets for Extended Logic Programs. *Artif. Intell.*, 109(1-2): 297–356.

Brewka, G.; Niemelä, I.; and Truszczynski, M. 2003. Answer Set Optimization. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence*, 867–872.

Cayrol, C.; Cohen, A.; and Lagasquie-Schiex, M. 2021. Higher-Order Interactions (Bipolar or not) in Abstract Argumentation: A State of the Art. In *Handbook of Formal Argumentation*, volume 2, chapter 1. College Publications.

Coste-Marquis, S.; Konieczny, S.; Marquis, P.; and Ouali, M. A. 2012. Weighted Attacks in Argumentation Frameworks. In *Proceedings of the Thirteenth International Conference on Principles of Knowledge Representation and Reasoning*.

Cyras, K. 2016. Argumentation-Based Reasoning with Preferences. In *Highlights of Practical Applications of Scalable Multi-Agent Systems. The PAAMS Collection - International Workshops of PAAMS*, 199–210.

Cyras, K.; Fan, X.; Schulz, C.; and Toni, F. 2018. Assumption-based Argumentation: Disputes, Explanations, Preferences. In *Handbook of Formal Argumentation*, volume 1, chapter 7. College Publications.

Delgrande, J. P.; Schaub, T.; and Tompits, H. 2003. A Framework for Compiling Preferences in Logic Programs. *Theory Pract. Log. Program.*, 3(2): 129–187.

Dung, P. M. 1995. On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games. *Artif. Intell.*, 77: 321–358.

Dung, P. M.; Thang, P. M.; and Son, T. C. 2019. On Structured Argumentation with Conditional Preferences. In *Proceeding of the Thirty-Third AAAI Conference on Artificial Intelligence*, 2792–2800.

Dunne, P. E.; and Bench-Capon, T. J. M. 2004. Complexity in Value-Based Argument Systems. In *Proceedings of the 14th European Conference on Logics in Artificial Intelligence*, 360–371.

Dunne, P. E.; Hunter, A.; McBurney, P.; Parsons, S.; and Wooldridge, M. J. 2011. Weighted argument systems: Basic definitions, algorithms, and complexity results. *Artif. Intell.*, 175(2): 457–486.

Dvorák, W.; and Dunne, P. E. 2018. Computational Problems in Formal Argumentation and their Complexity. In *Handbook of Formal Argumentation*, volume 1, chapter 13. College Publications.

Eiter, T.; Faber, W.; Leone, N.; and Pfeifer, G. 2003. Computing preferred answer sets by meta-interpretation in answer set programming. *Theory Pract. Log. Program.*, 3(4-5): 463–498.

Fazzinga, B.; Flesca, S.; and Furfaro, F. 2020. Revisiting the Notion of Extension over Incomplete Abstract Argumentation Frameworks. In *Proc. of IJCAI*, 1712–1718.

Fazzinga, B.; Flesca, S.; and Parisi, F. 2015. On the Complexity of Probabilistic Abstract Argumentation Frameworks. *ACM Trans. Comput. Log.*, 16(3): 22:1–22:39.

Fishburn, P. C. 1972. Even-chance lotteries in social choice theory. *Theory and Decision*, 3(1): 18–40.

Gabbay, D.; Giacomin, M.; Simari, G. R.; and Thimm, M., eds. 2021. *Handbook of Formal Argumentation*, volume 2. College Publications.

Ganzer-Ripoll, J.; Criado, N.; Lopez-Sanchez, M.; Parsons, S.; and Rodriguez-Aguilar, J. A. 2019. Combining social choice theory and argumentation: Enabling collective decision making. *Group Decision and Negotiation*, 28(1): 127–173.

Garcia, A. J.; Prakken, H.; and Simari, G. R. 2020. A Comparative Study of Some Central Notions of ASPIC+ and DeLP. *Theory Pract. Log. Program.*, 20(3): 358–390.

Gärdenfors, P. 1979. On definitions of manipulation of social choice functions. *Aggregation and Revelation of Preferences. North-Holland*, 78.

Goldsmith, J.; Lang, J.; Truszczynski, M.; and Wilson, N. 2008. The Computational Complexity of Dominance and Consistency in CP-Nets. *J. Artif. Intell. Res.*, 33: 403–432.

Gottifredi, S.; Cohen, A.; García, A. J.; and Simari, G. R. 2018. Characterizing acceptability semantics of argumentation frameworks with recursive attack and support relations. *Artif. Intell.*, 262: 336–368.

Greco, S.; Trubitsyna, I.; and Zumpano, E. 2007. On the Semantics of Logic Programs with Preferences. *J. Artif. Intell. Res.*, 30: 501–523.

Heyninck, J.; and Strasser, C. 2021. A Comparative Study of Assumption-based Argumentative Approaches to Reasoning with Priorities. *Journal of Applied Logics - IfCoLog Journal*, 8(3): 737–808.

Kaci, S.; van der Torre, L. W. N.; Vesic, S.; and Villata, S. 2021. Preference in Abstract Argumentation. In *Handbook of Formal Argumentation*, volume 2, chapter 3. College Publications.

Kelly, J. S. 1977. Strategy-proofness and social choice functions without singlevaluedness. *Econometrica: Journal of the Econometric Society*, 439–446.

Lukasiewicz, T.; and Malizia, E. 2019. Complexity results for preference aggregation over (m)CP-nets: Pareto and majority voting. *Artif. Intell.*, 272: 101–142.

Lukasiewicz, T.; and Malizia, E. 2022. Complexity results for preference aggregation over (m)CP-nets: Max and rank voting. *Artif. Intell.*, 303: 103636.

Mailly, J.; and Rossit, J. 2020. Argument, I Choose You! Preferences and Ranking Semantics in Abstract Argumentation. In *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning*, 647–651.

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

Modgil, S.; and Prakken, H. 2013. A general account of argumentation with preferences. *Artif. Intell.*, 195: 361–397.

Modgil, S.; and Prakken, H. 2014. The $ASPIC^+$ framework for structured argumentation: a tutorial. *Argument Comput.*, 5(1): 31–62.

Niskanen, A.; and Järvisalo, M. 2020. Algorithms for Dynamic Argumentation Frameworks: An Incremental SAT-Based Approach. In *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI)*, 849–856.

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

Prakken, H. 2010. An abstract framework for argumentation with structured arguments. *Argument & Computation*, 1(2): 93–124.

Prakken, H.; and Sartor, G. 1997. Argument-Based Extended Logic Programming with Defeasible Priorities. *J. Appl. Non Class. Logics*, 7(1): 25–75.

Rossi, F.; Venable, K. B.; and Walsh, T. 2004. mCP Nets: Representing and Reasoning with Preferences of Multiple Agents. In *Proceedings of the Nineteenth AAAI Conference on Artificial Intelligence*, 729–734.

Sakama, C.; and Inoue, K. 2000. Prioritized logic programming and its application to commonsense reasoning. *Artif. Intell.*, 123(1-2).

Sakama, C.; and Son, T. C. 2020. Epistemic Argumentation Framework: Theory and Computation. *J. Artif. Intell. Res.*, 69: 1103–1126.

Schaub, T.; and Wang, K. 2001. A Comparative Study of Logic Programs with Preference. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence*, 597–602.

Silva, R.; Sá, S.; and Alcântara, J. F. L. 2020. Semantics Hierarchy in Preference-Based Argumentation Frameworks. In *Proceedings of the 8th International Conference on Computational Models of Argument*, 339–346.

Teze, J. C.; Gottifredi, S.; García, A. J.; and Simari, G. R. 2020. An approach to generalizing the handling of preferences in argumentation-based decision-making systems. *Knowl. Based Syst.*, 189.

Villata, S.; Boella, G.; Gabbay, D. M.; and van der Torre, L. W. N. 2012. Modelling defeasible and prioritized support in bipolar argumentation. *Ann. Math. Artif. Intell.*, 66(1-4).

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