Computational Models of Argument
F. Toni et al. (Eds.)

doi:10.3233/FAIA220144

Abstract Argumentation with Conditional Preferences

Michael BERNREITER [1], Wolfgang DVOŘÁK and Stefan WOLTRAN
Institute of Logic and Computation, TU Wien, Austria

Abstract. In this paper, we study conditional preferences in abstract argumentation by introducing a new generalization of Dung-style argumentation frameworks (AFs) called Conditional Preference-based AFs (CPAFs). Each subset of arguments in a CPAF can be associated with its own preference relation. This generalizes existing approaches for preference-handling in abstract argumentation, and allows us to reason about conditional preferences in a general way. We conduct a principle-based analysis of CPAFs and compare them to related generalizations of AFs. Specifically, we highlight similarities and differences to Modgil's Extended AFs and show that our formalism can capture Value-based AFs.

Keywords. Abstract argumentation, conditional preferences, principles.

1. Introduction

Preferences in argumentation have been studied from various points of view, be it in terms of argument strength [1,2,3,4,5] or preferences between values [6,7]. Despite this, *conditional preferences* have received only limited attention in the field of argumentation. Dung et al. investigated conditional preferences in the setting of structured argumentation [8]. There, argumentation frameworks (AFs) are built from defeasible knowledge bases containing preference rules of the form $a_1,\ldots,a_n \to d_0 \succ d_1$, where d_0 and d_1 are defeasible rules. However, no work that deals with conditional preferences on the abstract level is known to us. This is in contrast to unconditional preferences, which are studied both in structured [9,10,11] and abstract [5,7] argumentation in the literature.

Conditional preferences can appear in many situations and formalisms. Dung et al. [8] demonstrate this with the help of an example, which we now adapt:[2]

Example 1. *Sherlock Holmes is investigating a murder. There are two suspects, Person 1 and Person 2. After analyzing the crime scene, Sherlock is sure:*

- I_1*: Person 1 or Person 2 is the culprit, but not both.*

Moreover, Sherlock adheres to the following rules:

- R_1*: If Person i has a motive but Person j, with $j \neq i$, does not, then this supports the case that Person i is the culprit.*

[1] Corresponding Author: Michael Bernreiter; E-mail: michael.bernreiter@tuwien.ac.at.
[2] We specify the example in natural language. See [8] for how Dung's original example can be modeled as a defeasible knowledge base with conditional preferences. Our example can be formalized similarly.

- *R_2: If Person i has an alibi but Person j, with $j \neq i$, does not, then this supports the case that Person j is the culprit.*
- *R_3: Alibis have more importance than motives.*

After interrogating the suspects, Sherlock concludes that:

- *C_1: Person 1 has a motive but Person 2 does not.*
- *C_2: Person 1 has an alibi but Person 2 does not.*

If C_1 is trusted, but C_2 is not, then this supports that Person 1 is the culprit. If C_2 is trusted then this supports that Person 2 is the culprit, regardless of our stance on C_1.

This example demonstrates the importance of conditional preferences in common reasoning tasks. We believe it is valuable to capture conditional preferences in argumentation not only on the structured level as Dung et al. [8] did, but also on the abstract level. Doing so will generalize existing formalisms for unconditional preferences in abstract argumentation and provide a more direct target formalism for structured approaches.

To this end, we introduce Conditional Preference-based AFs (CPAFs), where each subset of arguments S can be associated with its own preference relation $\succ_S$. Preferences are then resolved via so-called preference-reductions [5], which modify the attack relation based on the given preferences. As a consequence, S must be justified in view of its own preferences, i.e., S must be an extension in view of $\succ_S$.

We show that CPAFs generalize Preference-based AFs (PAFs), and demonstrate that they are capable of dealing with conditional preferences in a general manner. Moreover, we conduct a principle-based analysis of CPAF-semantics and show that especially complete and stable semantics preserve desirable properties of regular PAFs. Lastly, we compare CPAFs to related formalisms. Specifically, we show that CPAFs can capture other generalizations of AFs such as Value-based AFs (VAFs) [6,7] in a straightforward way, and compare CPAFs to Extended Argumentation Frameworks (EAFs) [12,13,14] in order to highlight similarities and differences.

2. Preliminaries

We first define (abstract) argumentation frameworks [15].

Definition 1. *An argumentation framework (AF) is a tuple $F = (A, R)$ where A is a finite set of arguments and $R \subseteq A \times A$ is an attack relation between arguments. Let $S \subseteq A$. We say S* attacks *b (in F) if $(a,b) \in R$ for some $a \in S$; $S_F^+ = \{b \in A \mid \exists a \in S : (a,b) \in R\}$ denotes the set of arguments attacked by S. An argument $a \in A$ is* defended *(in F) by S if $b \in S_F^+$ for each b with $(b,a) \in R$.*

Semantics for AFs are defined as functions σ which assign to each AF $F = (A, R)$ a set $\sigma(F) \subseteq 2^A$ of extensions [16]. We consider for σ the functions *cf* (conflict-free), *adm* (admissible), *com* (complete), *stb* (stable), *grd* (grounded), and *prf* (preferred).

Definition 2. *Let $F = (A, R)$ be an AF. A set $S \subseteq A$ is conflict-free (in F), written as $S \in cf(F)$, if there are no $a, b \in S$, such that $(a,b) \in R$. For $S \in cf(F)$ it holds that*

- *$S \in adm(F)$ if each $a \in S$ is defended by S in F;*
- *$S \in com(F)$ if $S \in adm(F)$ and each $a \in A$ defended by S in F is contained in S;*

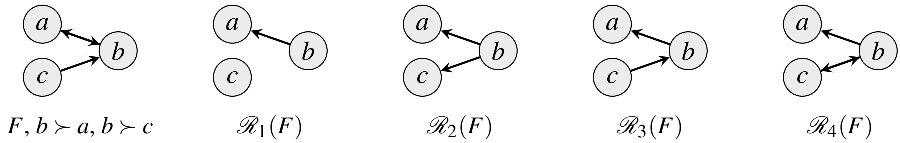

Figure 1. PAF F and its preference reducts from Example 2.

- $S \in stb(F)$ *if each* $a \in A \setminus S$ *is attacked by* S *in* F*;*
- $S \in grd(F)$ *if* $S \in com(F)$ *and there is no* $T \in com(F)$ *with* $T \subset S$*;*
- $S \in prf(F)$ *if* $S \in adm(F)$ *and there is no* $T \in adm(F)$ *with* $S \subset T$*;*

Preference-based AFs enrich regular AFs with preferences between arguments.

Definition 3. *A preference-based AF (PAF) is a triple* $F = (A, R, \succ)$ *where* (A, R) *is an AF and* $\succ$ *is an irreflexive and asymmetric binary relation over* A*.*

If a and b are arguments and $a \succ b$ holds then we say that a is stronger than b. An established method of resolving preferences in PAFs are so-called preference reductions, of which there exist four in the literature [5]. If in a PAF $(A, R, \succ)$ there is an attack $(a, b) \in R$ and a preference $b \succ a$ then (a, b) is called a critical attack. In other words, critical attacks are from weak to strong arguments. The preference-reductions resolve preferences by dealing with these critical attacks, e.g., by removing or reverting them.

Definition 4. *Given a PAF* $F = (A, R, \succ)$*, a corresponding AF* $\mathscr{R}_i(F) = (A, R')$ *is constructed via Reduction i, where* $i \in \{1, 2, 3, 4\}$*, as follows:*

- $i = 1$: $\forall a, b \in A: \ (a, b) \in R' \Leftrightarrow (a, b) \in R, b \not\succ a$
- $i = 2$: $\forall a, b \in A: \ (a, b) \in R' \Leftrightarrow ((a, b) \in R, b \not\succ a)$ *or* $((b, a) \in R, (a, b) \notin R, a \succ b)$
- $i = 3$: $\forall a, b \in A: \ (a, b) \in R' \Leftrightarrow ((a, b) \in R, b \not\succ a)$ *or* $((a, b) \in R, (b, a) \notin R)$
- $i = 4$: $\forall a, b \in A: \ (a, b) \in R' \Leftrightarrow ((a, b) \in R, b \not\succ a)$ *or* $((b, a) \in R, (a, b) \notin R, a \succ b)$ *or* $((a, b) \in R, (b, a) \notin R)$

The preference-based variant of a semantics σ *relative to Reduction i is defined as* $\sigma_p^i(F) = \sigma(\mathscr{R}_i(F))$*.*

Intuitively, Reduction 1 removes critical attacks while Reduction 2 reverts them. Reduction 3 removes critical attacks, but only if the stronger argument also attacks the weaker one. Reduction 4 can be seen as a combination of Reduction 2 and 3. Note that on symmetric attacks, all four reductions function in the same way. The following example demonstrates the reductions and PAF-semantics.

Example 2. *Consider the PAF* $F = (\{a, b, c\}, \{(a, b), (b, a), (c, b)\}, \succ)$ *with* $b \succ a$ *and* $b \succ c$*. Figure 1 depicts* F *as well as* $\mathscr{R}_i(F)$*,* $i \in \{1, 2, 3, 4\}$*. It can be checked that, for Reduction 1,* $adm_p^1(F) = adm(\mathscr{R}_i(F)) = \{\emptyset, \{b\}, \{c\}, \{b, c\}\}$ *and therefore* $com_p^1(F) = prf_p^1(F) = stb_p^1(F) = \{\{b, c\}\}$*. If we use Reduction 2 for example we get different extensions, namely* $adm_p^2(F) = \{\emptyset, \{b\}\}$ *and* $com_p^2(F) = prf_p^2(F) = stb_p^2(F) = \{\{b\}\}$*.*

A principle-based analysis of the four preference reductions was conducted for complete, grounded, preferred, and stable semantics [4,5]. To this end, the following six PAF-properties were laid out and investigated.

	$\mathscr{R}_1$	$\mathscr{R}_2$	$\mathscr{R}_3$	$\mathscr{R}_4$
$P1$	$\times$	$CGPS$	$CGPS$	$CGPS$
$P2$	$\times$	$\times$	CS	$\times$
$P3$	$\times$	$\times$	CS	$\times$
$P4$	$\times$	$\times$	CGS	$\times$
$P5$	$\times$	$\times$	CG	$\times$
$P6$	G	G	$CGPS$	G

Table 1. Satisfaction of various PAF-principles. C stands for complete, G for grounded, P for preferred, and S for stable. $\times$ indicates that none of those four semantics satisfy this principle.

Definition 5. *Let σ_p^i be a PAF-semantics. Let $\succ' \subseteq (A \times A)$ be irreflexive and asymmetric.*

- *σ_p^i satisfies $P1$ (conflict-freeness) iff for all PAFs $F = (A, R, \succ)$ there is no $S \in \sigma_p^i(F)$ such that $\{a, b\} \subseteq S$ and $(a, b) \in R$.*
- *σ_p^i satisfies $P2$ (preference selects extensions 1) iff $\sigma_p^i(A, R, \succ \cup \succ') \subseteq \sigma_p^i(A, R, \succ)$ holds for all PAFs $(A, R, \succ)$ and all $\succ'$.*
- *σ_p^i satisfies $P3$ (preference selects extensions 2) iff $\sigma_p^i(A, R, \succ) \subseteq \sigma_p^i(A, R, \emptyset)$ holds for all PAFs $(A, R, \succ)$.*
- *σ_p^i satisfies $P4$ (extension refinement) iff for all $S' \in \sigma_p^i(A, R, \succ \cup \succ')$ there exists some $S \in \sigma_p^i(A, R, \succ)$ such that $S \subseteq S'$.*
- *σ_p^i satisfies $P5$ (extension growth) iff $\bigcap(\sigma_p^i(A, R, \succ)) \subseteq \bigcap(\sigma_p^i(A, R, \succ \cup \succ'))$ holds for all PAFs $(A, R, \succ)$ and all $\succ'$.*
- *σ_p^i satisfies $P6$ (number of extensions) iff $|\sigma_p^i(A, R, \succ \cup \succ')| \leq |\sigma_p^i(A, R, \succ)|$ holds for all PAFs $(A, R, \succ)$ and all $\succ'$.*

Intuitively, $P1$ states that if there is an attack between two arguments, then there is no extension containing both of them. $P2$ expresses that adding more preferences to a PAF can exclude extensions, but not introduce them. $P3$ is a special case of $P2$. $P4$ states that adding preferences means extensions will be supersets of extensions in the original PAF. $P5$ says that adding preferences will preserve skeptically accepted arguments, and might cause new arguments to be skeptically accepted. $P6$ expresses that the number of extensions will not grow if new preferences are added.

Table 1 shows which semantics satisfy which principle. In addition to these fundamental principles [4], four more principles were introduced later [5], but we do not consider them at this point and leave them for future work.

3. Conditional Preference-Based Argumentation Frameworks

As argued in the introduction, our aim is to provide a framework for reasoning about conditional preferences in abstract argumentation. This means that arguments themselves must be capable of expressing preferences, and that those argument-bound preferences are relevant only if the corresponding arguments are themselves accepted. How this is implemented must be considered carefully, as Example 1 demonstrates. There, the fact that Person 1 has a motive (let us refer to this as m_1) and the fact that Person 1 has an alibi (a_1) result in opposing preferences. When accepting both m_1 and a_1 it seems natural to combine these opposing preferences, i.e., to cancel them. But this does not allow us to express that alibis are more important than motives, as required in Example 1. Therefore,

F, $c_1 \succ_S c_2$ if $m_1 \in S$, $a_1 \notin S$ and $c_2 \succ_S c_1$ if $a_1 \in S$ $\quad$ $\mathscr{R}_i^S(F)$, $m_1 \in S$, $a_1 \notin S$ $\quad$ $\mathscr{R}_i^S(F)$, $a_1 \in S$

Figure 2. CPAF F and its preference-reducts from Example 3.

we need to define our formalism in a general way such that the joint acceptance of arguments must not necessarily result in the combination of their associated preferences. We solve this by mapping each subset S of arguments to a separate preference relation $\succ_S$.

Definition 6. *A Conditional PAF (CPAF) is a triple $F = (A, R, cond)$, where (A, R) is an AF and $cond\colon 2^A \to 2^{(A \times A)}$ is a function that maps each set of arguments $S \subseteq A$ to an irreflexive and asymmetric binary relation $\succ_S$ over A.*

Note that we set no restriction on how exactly conditional preferences are represented. This is deliberate, as we wish to stay as general as possible. In practice, succinct representations could be achieved, e.g., by expressing the *cond*-function via rules of the form $F \to x \succ y$ where F is a propositional formula over the arguments.

Just as in regular PAFs, preferences in CPAFs are resolved with the help of the four preference-reductions (cf. Definition 4). A set of arguments S is an extension of some CPAF if it is an extension relative to its associated preference relation $cond(S)$.

Definition 7. *Let $F = (A, R, cond)$ be a CPAF and let $S \subseteq A$. The S-reduct of F with respect to a preference reduction $i \in \{1,2,3,4\}$ is defined as $\mathscr{R}_i^S(F) = \mathscr{R}_i(A, R, cond(S))$. Given an AF semantics σ, $S \in \sigma_{cp}^i(F)$ iff $S \in \sigma(\mathscr{R}_i^S(F))$.*

Using CPAFs we can easily formalize our Sherlock Holmes example.

Example 3. *We continue Example 1 and introduce two arguments c_1 and c_2 expressing that Person 1 (resp. Person 2) is the culprit. Moreover, we introduce m_1 and a_1 to express that Person 1 has a motive (resp. alibi) but Person 2 does not. c_1 and c_2 attack each other while m_1 and a_1 have no incoming or outgoing attacks, but rather express preferences. Formally, we model this via the CPAF $F = (\{c_1, c_2, m_1, a_1\}, \{(c_1, c_2), (c_2, c_1)\}, cond)$ with cond such that $c_1 \succ_S c_2$ if $m_1 \in S$ but $a_1 \notin S$, $c_2 \succ_S c_1$ if $a_1 \in S$, and $cond(S) = \emptyset$ for all other $S \subseteq A$. Figure 2 depicts F and the S-reducts of F. Note that m_1 and a_1 are unattacked in all S-reducts of F. Therefore, both arguments must be part of any σ_{cp}^i-extension for $\sigma \in \{grd, com, prf, stb\}$ and we can conclude that $\sigma_{cp}^i(F) = \{\{m_1, a_1, c_2\}\}$.*

Note that the preferred semantics defined above do not maximize over all admissible sets of a CPAF, but rather over all admissible sets in the given S-reduct. This means that if there is a set S that is admissible in the S-reduct of F, but there is also some $T \supset S$ that is admissible in the S-reduct of F, then S is not preferred in F. But this T does not have to be admissible in F, since it might not be admissible in the T-reduct of F. Thus, the following alternative semantics may be considered more natural:

Definition 8. *Let $F = (A, R, cond)$ be a CPAF and let $S \subseteq A$. Then $S \in prf\text{-}glb_{cp}^i(F)$ iff $S \in adm_{cp}^i(F)$ and there is no T such that $S \subset T$ and $T \in adm_{cp}^i(F)$.*

Intuitively, $prf\text{-}glb^i_{cp}$ maximizes *globally* over all admissible sets of a CPAF, while prf^i_{cp} maximizes *locally* over the admissible sets of the given S-reduct.

Example 4. *Let F be the CPAF from Example 3 and recall that $prf^i_{cp}(F) = \{\{m_1, a_1, c_2\}\}$. Observe that $\{m_1, c_1\}$ is not preferred in the $\{m_1, c_1\}$-reduct of F, but it is a subset-maximal admissible set in F. Thus, $prf\text{-}glb^i_{cp}(F) = \{\{m_1, a_1, c_2\}, \{m_1, c_1\}\}$.*

The difference between the two variants is not only philosophical, but impacts fundamental properties for maximization-based semantics such as I-maximality. A semantics σ^i_{cp} is I-maximal if, for all CPAFs F and all $S, T \in \sigma^i_{cp}(F)$, $S \subseteq T$ implies $S = T$.

Proposition 1. *$prf\text{-}glb^i_{cp}$ is I-maximal, but prf^i_{cp} is not, where $i \in \{1,2,3,4\}$.*

Proof. I-maximality of $prf\text{-}glb^i_{cp}$ follows from Definition 8. For our counter examples we consider the preference-reductions separately. Reduction 1: consider the CPAF $F = (\{a,b\}, \{(a,b)\}, cond\}$ with $cond$ such that $b \succ_{\{a,b\}} a$. Then $\{a\} \in prf^1_{cp}(F)$ and $\{a,b\} \in prf^1_{cp}(F)$. Reductions 2 and 4: consider the CPAF $F' = (\{a,b,c\}, \{(a,b),(b,c),(c,a)\}, cond\}$ with $cond$ such that $a \succ_{\{a\}} c$. Then $\emptyset \in prf^i_{cp}(F')$ and $\{a\} \in prf^i_{cp}(F')$. Reduction 3: consider the CPAF $F'' = (\{a,b,c\}, \{(a,b),(b,a),(b,c),(c,a)\}, cond\}$ with $cond$ such that $a \succ_\emptyset b$. Then $\emptyset \in prf^3_{cp}(F'')$ and $\{b\} \in prf^3_{cp}(F'')$. □

One may be tempted to deduce from the above proposition that $prf\text{-}glb^i_{cp}$ is more suitable as a default preferred semantics than prf^i_{cp}. However, we will see in Section 5.1 that prf^i_{cp} allows us to capture the problems of subjective/objective acceptance in VAFs in a natural way. In our subsequent analysis of CPAFs we consider both prf^i_{cp} and $prf\text{-}glb^i_{cp}$. Like preferred semantics, stable semantics satisfy I-maximality on regular AFs. Interestingly, on CPAFs, this depends on the preference-reduction.

Proposition 2. *stb^1_{cp} is not I-maximal, but stb^j_{cp} is, where $j \in \{2,3,4\}$.*

Proof. For stb^1_{cp} we can use the same counter-example as for prf^1_{cp} (cf. Proposition 1). For stb^j_{cp} with $j \in \{2,3,4\}$ we proceed by contradiction: assume there is a CPAF $F = (A, R, cond)$ with $S, T \in stb^j_{cp}(F)$ such that $S \subset T$. Then there is $x \in T$ such that $x \notin S$. Since $S \in stb^j_{cp}(F)$ there is $y \in S$ such that $(y,x) \in \mathscr{R}^S_j(F)$. Reductions 2, 3, and 4 do not remove conflicts between arguments, and thus either $(y,x) \in R$ or $(x,y) \in R$. Therefore, $(y,x) \in \mathscr{R}^T_j(F)$ or $(x,y) \in \mathscr{R}^T_j(F)$. But $y \in S$ implies $y \in T$, i.e., $T \notin cf^j_{cp}(F)$. □

Another interesting point is that grounded extensions are not necessarily unique in CPAFs: consider $F = (\{a,b\}, \{(a,b)\}, cond)$ with $cond$ such that $b \succ_{\{b\}} a$. Then $\{a\} \in grd^2_{cp}(F)$ and $\{b\} \in grd^2_{cp}(F)$. We stress that each grounded extension S is still unique in the S-reduct of the given CPAF and thus unique with respect to its own preferences.

Lastly, by the following proposition we express that all CPAF-semantics considered here generalize their corresponding PAF-semantics, i.e., that CPAFs generalize PAFs.

Proposition 3. *Let $F = (A, R, cond)$ be a CPAF such that the preference function $cond$ maps every set of arguments to the same binary relation, i.e., there is some $\succ$ such that $cond(S) = \succ$ for all $S \subseteq A$. Let $\sigma \in \{cf, adm, stb, com, prf, grd\}$. Then $\sigma^i_{cp}(F) = \sigma^i_p(A, R, \succ)$. Furthermore, $prf\text{-}glb^i_{cp}(F) = prf^i_p(A, R, \succ)$.*

4. Principle-Based Analysis

In this section, we generalize the principles of Kaci et al. for PAFs (cf. Definition 5) to account for conditional preferences. We then investigate by which semantics these principles are satisfied, and show that there are differences to the case of regular PAFs.

Definition 9. *Let σ_{cp}^i be a CPAF-semantics. In the following, given a CPAF $(A,R,cond)$, we denote by $cond'$ an arbitrary function such that $cond(S) \subseteq cond'(S)$ for all $S \subseteq A$. Furthermore, $cond_\emptyset(S) = \emptyset$ for all $S \subseteq A$.*

- *σ_{cp}^i satisfies $P1^*$ (conflict-freeness) iff for all CPAFs $F = (A,R,cond)$ there is no $S \in \sigma_{cp}^i(F)$ such that $\{a,b\} \subseteq S$ and $(a,b) \in R$.*
- *σ_{cp}^i satisfies $P2^*$ (preference selects extensions) iff for all CPAFs $(A,R,cond)$ it holds that $\sigma_{cp}^i(A,R,cond') \subseteq \sigma_{cp}^i(A,R,cond)$.*
- *σ_{cp}^i satisfies $P3^*$ (preference selects extensions 2) iff for all CPAFs $(A,R,cond)$ it holds that $\sigma_{cp}^i(A,R,cond) \subseteq \sigma_{cp}^i(A,R,cond_\emptyset)$.*
- *σ_{cp}^i satisfies $P4^*$ (extension refinement) iff for all $S' \in \sigma_{cp}^i(A,R,cond')$ there exists some $S \in \sigma_{cp}^i(A,R,cond)$ such that $S \subseteq S'$.*
- *σ_{cp}^i satisfies $P5^*$ (extension growth) iff for all CPAFs $(A,R,cond)$ it holds that $\bigcap(\sigma_{cp}^i(A,R,cond)) \subseteq \bigcap(\sigma_{cp}^i(A,R,cond'))$.*
- *σ_{cp}^i satisfies $P6^*$ (number of extensions) iff for all CPAFs $(A,R,cond)$ it holds that $|\sigma_{cp}^i(A,R,cond')| \leq |\sigma_{cp}^i(A,R,cond)|$.*

The following lemma establishes some relationships between the CPAF-principles and is a generalization of known relationships between PAF-principles [5].

Lemma 4. *If σ_{cp}^i satisfies $P2^*$ then it also satisfies $P3^*$, $P4^*$, and $P6^*$. If σ_{cp}^i always returns at least one extension, and if it satisfies $P2^*$, then it also satisfies $P5^*$.*

Observe that, since CPAFs are a generalization of PAFs (cf. Proposition 3), a CPAF-semantics σ_{cp}^i can not satisfy Pj^* if the corresponding PAF-semantics σ_p^i does not satisfy Pj. Moreover, it is obvious that $P1^*$ is still satisfied under Reductions 2, 3, and 4, as conflicts are not removed by these reductions even if we consider conditional preferences. We can also show that satisfaction of $P2$ carries over from PAFs to CPAFs.

Lemma 5. *If some σ_p^i satisfies P2 then σ_{cp}^i satisfies $P2^*$.*

Proof. Assume σ_{cp}^i does not satisfy $P2^*$. Then there is a CPAF $F = (A,R,cond)$ and $cond'$ with $cond(S) \subseteq cond'(S)$ for all $S \subseteq A$ such that $\sigma_{cp}^i(A,R,cond') \not\subseteq \sigma_{cp}^i(A,R,cond)$. Thus, there is $E \subseteq A$ such that $E \in \sigma_{cp}^i(A,R,cond')$ but $E \notin \sigma_{cp}^i(A,R,cond)$. Then $E \in \sigma(\mathscr{R}_i(A,R,cond'(E)))$ but $E \notin \sigma(\mathscr{R}_i(A,R,cond(E)))$, i.e., σ_p^i does not satisfy $P2$. □

Lemma 5 implies that complete and stable semantics satisfy $P2^*$ on CPAFs under Reduction 3. By Lemma 4 these semantics also satisfy $P3^*$, $P4^*$, and $P6^*$. However, we can not use Lemma 4 to show that complete semantics satisfy $P5^*$, since conditional preferences allow for frameworks without complete extensions. Indeed, we can find a counter-example in this case. Counter-examples for the satisfaction of various principles can also be found for grounded semantics and both variants of the preferred semantics.

	$\mathscr{R}_1$	$\mathscr{R}_2$	$\mathscr{R}_3$	$\mathscr{R}_4$
$P1^*$	$\times$	*CGPS*	*CGPS*	*CGPS*
$P2^*$	$\times$	$\times$	*CS*	$\times$
$P3^*$	$\times$	$\times$	*CS*	$\times$
$P4^*$	$\times$	$\times$	*CS*	$\times$
$P5^*$	$\times$	$\times$	$\times$	$\times$
$P6^*$	$\times$	$\times$	*CS*	$\times$

Table 2. Satisfaction of CPAF-principles. C stands for complete, G for grounded, P for preferred (local and global maximization), and S for stable. $\times$ indicates that none of those semantics satisfy this principle.

Lemma 6. *com^3_{cp} does not satisfy $P5^*$. grd^i_{cp}, with $i \in \{1,2,3,4\}$, does not satisfy any of $P4^*$, $P5^*$, or $P6^*$. Moreover, prf^3_{cp} and $prf\text{-}glb^3_{cp}$ do not satisfy $P6^*$.*

Proof. For complete semantics, consider $A = \{a,b\}$, $R = \{(a,b),(b,a)\}$, *cond* such that $a \succ_\emptyset b$ and $a \succ_{\{b\}} b$, as well as $cond'$ such that $a \succ'_\emptyset b$, $a \succ'_{\{b\}} b$, and $b \succ'_{\{a\}} a$. Then $com^3_{cp}(A,R,cond) = \{\{a\}\}$ while $com^3_{cp}(A,R,cond') = \emptyset$.

For grounded semantics, consider $A = \{a,b\}$, $R = \{(a,b),(b,a)\}$, *cond* such that $a \succ_\emptyset b$ and $a \succ_{\{a\}} b$, as well as $cond'$ such that $a \succ'_\emptyset b$, $a \succ'_{\{a\}} b$, and $b \succ'_{\{b\}} a$. Then $grd^i_{cp}(A,R,cond) = \{\{a\}\}$ while $grd^i_{cp}(A,R,cond') = \{\{a\},\{b\}\}$.

For preferred semantics, consider $A = \{a,b,c\}$, $R = \{(a,c),(c,a),(b,c),(c,b),(c,c)\}$, *cond* such that $cond(S) = \emptyset$ for all $S \subseteq A$, and $cond'$ such that $c \succ'_{\{b\}} a$, $c \succ'_{\{a\}} b$, $c \succ'_{\{a,b\}} a$, and $c \succ'_{\{a,b\}} b$. Then $prf^3_{cp}(A,R,cond) = prf\text{-}glb^3_{cp}(A,R,cond) = \{\{a,b\}\}$ while $prf^3_{cp}(A,R,cond') = prf\text{-}glb^3_{cp}(A,R,cond') = \{\{a\},\{b\}\}$. □

The above results constitute an exhaustive investigation of the six CPAF-principles for all semantics considered in this paper. Thus, we can conclude:

Theorem 7. *The satisfaction of CPAF-principles depicted in Table 2 holds.*

To summarize, complete and stable semantics preserve the satisfaction of PAF-principles in almost all cases. Grounded semantics no longer satisfy any of the principles on CPAFs except $P1^*$ (conflict-freeness) since grounded extensions are not unique on CPAFs, and since there are even CPAFs without a grounded extension (cf. Lemma 6). Unlike on PAFs, complete semantics do not satisfy $P5^*$ (extension growth) under Reduction 3. Furthermore, neither variant of preferred semantics satisfies $P6^*$ (number of extensions) under Reduction 3.

5. Related Formalisms

We now investigate the connection between CPAFs and two related formalisms. First, we show that Value-based Argumentation Frameworks (VAFs) [6,7] can be captured by CPAFs in a straightforward way. Secondly, we consider Extended Argumentation Frameworks (EAFs) [12,13,14] and highlight similarities and differences to CPAFs.

5.1. Capturing Value-Based Argumentation

VAFs, similarly to CPAFs, are capable of dealing with multiple preference relations. But, in contrast to CPAFs, these preferences are not over individual arguments but over values

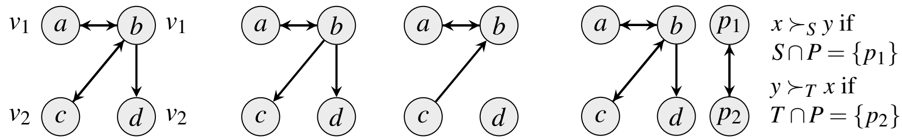

(a) Initial VAF (b) AF for p_1 (c) AF for p_2 (d) CPAF ($x \in \{a,b\}, y \in \{c,d\}$)

Figure 3. A VAF with two audiences p_1 ($v_1 \succ v_2$) and p_2 ($v_2 \succ v_1$) translated to a CPAF.

associated with arguments. Which values are preferred depends on the audience. A set of arguments may then be accepted in view of one audience, but not in view of another.

More formally, a VAF is a quintuple (A,R,V,val,P) such that (A,R) is an AF, V is a set of values, $val\colon A \to V$ is a mapping from arguments to values, and P is a finite set of audiences. Each audience $p \in P$ is associated with a preference relation $\succ_p$ over values, and $F_P = (A,R,V,val,\succ_p)$ is called an audience-specific VAF (AVAF). The extensions of VAFs are determined for each audience separately. Specifically, an argument x successfully attacks y in F_p iff $(x,y) \in R$ and $val(y) \not\succ_p val(x)$. Conflict-freeness and admissibility are then defined over these successful attacks. In essence, this boils down to using Reduction 1 on F_p, i.e., deleting attacks that contradict the preference ordering.

For example, Figure 3a shows a VAF with two values v_1 and v_2. Let us say there are two audiences in this VAF, p_1 with the preference $v_1 \succ v_2$ and p_2 with $v_2 \succ v_1$. The AFs associated with p_1 and p_2, i.e., the AFs containing only the successful attacks in the AVAFs of p_1 and p_2, are depicted in Figures 3b and 3c.

The reasoning tasks typically associated with VAFs are those of subjective and objective acceptance. Let $F = (A,R,V,val,P)$ be a VAF and $x \in A$. Then x is subjectively accepted in F iff there is $p \in P$ such that x is in a preferred extension of the AVAF $(A,R,V,val,\succ_p)$. Similarly, x is objectively accepted in F iff for all $p \in P$ we have that x is in all preferred extensions of the AVAF $(A,R,V,val,\succ_p)$.

We now provide a translation where an arbitrary VAF $F = (A,R,V,val,P)$ is transformed into a CPAF $Tr(F) = (A',R',cond)$ such that the subjectively and objectively accepted arguments in F correspond to the credulously and skeptically preferred arguments[3] in $Tr(F)$ respectively. Firstly, each audience in the initial VAF becomes an argument in our CPAF, i.e., $A' = A \cup P$. Secondly, the attacks of the VAF are preserved and symmetric attacks are added between all audience-arguments, i.e., $R' = R \cup \{(p,p'),(p',p) \mid p,p' \in P\}$. Lastly, the preferences in our CPAF correspond to the preferences of each audience and are controlled by the newly introduced audience-arguments, i.e., $cond$ is defined such that for $S \subseteq A'$ we have $a \succ_S b$ iff there is $p \in P$ with $S \cap P = \{p\}$ and $val(a) \succ_p val(b)$. See Figure 3 for an example of this transformation.

Observe that the successful attacks in some AVAF $F_p = (A,R,V,val,\succ_p)$ are also attacks in $\mathscr{R}_1^{S\cup\{p\}}(Tr(F))$, where $S \subseteq A$, and vice versa. Thus, the admissible sets in the initial VAF F stand in direct relationship to the admissible sets in our constructed CPAF.

Lemma 8. *Let $F = (A,R,V,val,P)$ be a VAF, $S \subseteq A$, and $p \in P$. Then S is admissible in the AVAF $F_p = (A,R,V,val,\succ_p)$ iff $S \cup \{p\} \in adm^1_{cp}(Tr(F))$.*

[3] As for regular AFs, we say that an argument x is credulously (resp. skeptically) preferred in a CPAF w.r.t. Reduction i iff $x \in S$ for some (resp. for all) $S \in prf^i_{cp}(F)$.

Note that all audience-arguments in $Tr(F)$ attack each other, i.e., an admissible set in $Tr(F)$ contains at most one audience-argument. In fact, each audience-argument defends itself, and thus every preferred extension in $Tr(F)$ must contain exactly one audience-argument $p \in P$ if we appeal to the prf^1_{cp}-semantics. Therefore, the direct correspondence between admissible sets observed in Lemma 8 carries over to preferred extensions.

Theorem 9. *Given a VAF $F = (A, R, V, val, P)$, $x \in A$ is subjectively (resp. objectively) accepted in F iff x is credulously (resp. skeptically) preferred in $Tr(F)$ w.r.t. Reduction* 1.

Our translation highlights the versatility of our formalism. On the one hand, conditional preferences can be tied to dedicated arguments (in this case the audience-arguments). On the other hand, these dedicated arguments themselves may be part of the argumentation process. Note that we used CPAFs with Reduction 1 since preferences in VAFs are usually handled by deleting attacks. However, our approach also allows for the use of other preference-reductions in VAF-settings.

5.2. Relationship to Extended Argumentation Frameworks

EAFs allow arguments to express preferences between other arguments by permitting attacks themselves to be attacked. While EAFs are related to our CPAFs conceptually, we will see that there are crucial differences in how exactly preferences are handled.

Formally, an EAF is a triple (A, R, D) such that (A, R) is an AF, $D \subseteq A \times R$, and if $(a, (b, c)), (a', (c, b)) \in D$ then $(a, a'), (a', a) \in R$. The definition of admissibility in EAFs is quite involved and requires so-called reinstatement sets. Essentially, a set of arguments S is admissible in an EAF if all arguments $x \in S$ are defended from other arguments $y \in A \setminus S$, and if all attacks (z, y) used for defending x are in turn defended from attacks on attacks $(w, (z, y))$ and thus reinstated. It is possible that a chain of such reinstatements is required which is formalized with the aforementioned reinstatement sets. Formally defining these concepts is not necessary for our purposes, but the corresponding definitions can be found in [12]. Observe that the notion of attacks on attacks in EAFs is similar to Reduction 1 in the sense that attacks between arguments can be unsuccessful, but they are never reversed. Therefore, we will compare EAFs to CPAFs with Reduction 1.

Recall our Sherlock Holmes example from the introduction (Example 1) that we modeled as a CPAF (Example 3). Let us first consider a slimmed-down variation without an argument stating that Person 1 has an alibi. We can model this as an EAF with three arguments c_1 (Person 1 is the culprit), c_2 (Person 2 is the culprit), and m_1 (Person 1 has a motive) in which m_1 attacks the attack from c_2 to c_1. The corresponding EAF is depicted in Figure 4b. Compare this to the formalization via a CPAF in Figure 4a. Note that $\{c_1\}$ is admissible in the EAF but $\{c_2\}$ is not since (c_2, c_1) is used to defend against (c_1, c_2) but not reinstated against $(m_1, (c_2, c_1))$. In the CPAF, $\{c_2\}$ is admissible (but not stable).

This simple example highlights a fundamental difference in how preferences are viewed in the two formalisms. In CPAFs, preferences are relevant exactly if the argument that expresses them (e.g. m_1) are part of the set under inspection. In EAFs, preference are relevant even if the argument that expresses them is not accepted. Modgil [12] states that admissibility for EAFs was defined in this way because it was deemed important to satisfy Dung's Fundamental Lemma [15], which says that if S is admissible and x is acceptable w.r.t. S then $S \cup \{x\}$ is admissible. This Fundamental Lemma is not satisfied in our CPAFs. However, in our opinion, this is no drawback but rather a necessary property

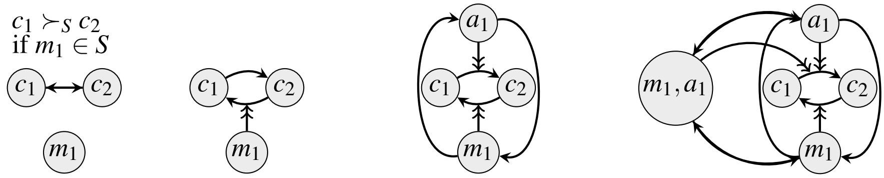

(a) Simple CPAF (b) Simple EAF (c) Conflicting preferences (d) Combining preferences

Figure 4. The Sherlock Holmes example modeled via EAFs and a simple CPAF.

of formalisms that can deal with conditional preferences in a flexible way. For example, in Figure 4a it is clear that $\{c_2\}$ should be admissible since, when considering only admissibility, we are not forced to include the unattacked m_1, i.e., we do not have to accept that Person 1 has a motive. The inclusion of unattacked arguments in CPAFs is handled via more restrictive approaches such as stable or preferred semantics, as usual.

Another difference between CPAFs and EAFs becomes clear when considering the entire Sherlock Holmes example. Recall our formalization for CPAFs (cf. Figure 2). In order to express our preference in case Person 1 has an alibi we extend our EAF from Figure 4b by adding an attack from a_1 to the attack (c_1, c_2), as shown in Figure 4c. Note that a_1 and m_1 must attack each other in this EAF by definition since they express conflicting preferences. But this formalization is unsatisfactory since it should be possible for Person 1 to have both a motive and an alibi. The fact that the preference of one argument may change in view of another argument must be modeled indirectly in EAFs. For example, we can introduce an additional argument to express that Person 1 has both a motive and an alibi. This is depicted in Figure 4d. Thus, we can see that CPAFs allow for more flexibility when combining preferences associated with several arguments.

To summarize, CPAFs are designed to express conditional preferences in abstract argumentation, whereas preferences in EAFs are unconditional in the sense that they may always influence the argumentation process, even if the argument associated with the preference is not accepted. Moreover, since our CPAFs can make use of all four preference reductions, they allow for more flexibility in how preferences are handled compared to EAFs, in which unsuccessful attacks are always deleted. However, the two formalisms are similar in that arguments are capable of reasoning about the argumentation process itself, i.e., they constitute a form of metalevel argumentation [17].

6. Conclusion

In this paper, we introduce Conditional Preference-based AFs (CPAFs) which generalize PAFs and allow to flexibly handle conditional preferences in abstract argumentation. We show that the satisfaction of I-maximality can depend on how maximization is dealt with (in case of preferred semantics) and on which preference-reduction is chosen (in case of stable semantics). We conduct a principle-based analysis for CPAFs and show that complete and stable semantics satisfy the same principles as on PAFs in most cases while grounded semantics no longer satisfy the majority of principles. Moreover, we compare CPAFs to related formalisms: on the one hand we show that CPAFs can be used to capture VAFs via a straightforward translation; on the other hand, we demonstrate that CPAFs exhibit significant differences to EAFs in terms of how preferences are handled.

For future work, we plan to introduce an alternative grounded semantics which enforces unique extensions, examine the computational complexity of CPAFs, and consider restricted (e.g. transitive or linear) preference orderings. Moreover, we intend to investigate the relationship between CPAFs and existing approaches in structured argumentation [8] in detail. Related to this last point, it may also be interesting to see whether conditional preferences can be adapted to other formalisms such as bipolar argumentation frameworks [18], in which both attack and support relations are present. As for preference representation, it could be investigated how existing formalisms designed to handle conditional preferences such as CP-nets [19] can be used in the context of CPAFs.

Acknowledgments We thank the anonymous reviewers for their valuable feedback. This work was funded by the Austrian Science Fund (FWF) under grant Y698 and by the Vienna Science and Technology Fund (WWTF) under grant ICT19-065.

References

[1] Amgoud L, Cayrol C. On the Acceptability of Arguments in Preference-based Argumentation. In: Proc. UAI'98. Morgan Kaufmann; 1998. p. 1-7.

[2] Amgoud L, Cayrol C. A Reasoning Model Based on the Production of Acceptable Arguments. Ann Math Artif Intell. 2002;34(1-3):197-215.

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

[4] Kaci S, van der Torre LWN, Villata S. Preference in Abstract Argumentation. In: Proc. COMMA'18. vol. 305 of FAIA. IOS Press; 2018. p. 405-12.

[5] Kaci S, van der Torre LWN, Vesic S, Villata S. Preference in Abstract Argumentation. In: Handbook of Formal Argumentation, Volume 2. College Publications; 2021. p. 211-48.

[6] Bench-Capon TJM, Doutre S, Dunne PE. Audiences in argumentation frameworks. Artif Intell. 2007;171(1):42-71.

[7] Atkinson K, Bench-Capon TJM. Value-based Argumentation. IfCoLog Journal of Logic and its Applications. 2021;8(6):1543-88.

[8] Dung PM, Thang PM, Son TC. On Structured Argumentation with Conditional Preferences. In: Proc. AAAI'19. AAAI Press; 2019. p. 2792-800.

[9] Modgil S, Prakken H. Reasoning about Preferences in Structured Extended Argumentation Frameworks. In: Proc. COMMA'10. vol. 216 of FAIA. IOS Press; 2010. p. 347-58.

[10] Modgil S, Prakken H. A general account of argumentation with preferences. Artif Intell. 2013;195:361-97.

[11] Modgil S, Prakken H. Abstract Rule-Based Argumentation. FLAP. 2017;4(8).

[12] Modgil S. Reasoning about Preferences in Argumentation Frameworks. Artif Intell. 2009;173(9-10):901-34.

[13] Dunne PE, Modgil S, Bench-Capon TJM. Computation in Extended Argumentation Frameworks. In: Proc. ECAI'10. vol. 215 of FAIA. IOS Press; 2010. p. 119-24.

[14] Baroni P, Cerutti F, Giacomin M, Guida G. Encompassing Attacks to Attacks in Abstract Argumentation Frameworks. In: Proc. ECSQARU'09. vol. 5590 of LNCS. Springer; 2009. p. 83-94.

[15] Dung PM. On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games. Artif Intell. 1995;77(2):321-58.

[16] Baroni P, Caminada M, Giacomin M. Abstract Argumentation Frameworks and Their Semantics. In: Handbook of Formal Argumentation. College Publications; 2018. p. 159-236.

[17] Modgil S, Bench-Capon TJM. Metalevel argumentation. J Log Comput. 2011;21(6):959-1003.

[18] Amgoud L, Cayrol C, Lagasquie-Schiex M, Livet P. On bipolarity in argumentation frameworks. Int J Intell Syst. 2008;23(10):1062-93.

[19] Boutilier C, Brafman RI, Domshlak C, Hoos HH, Poole D. CP-nets: A Tool for Representing and Reasoning with Conditional Ceteris Paribus Preference Statements. J Artif Intell Res. 2004;21:135-91.