

Abstract argumentation with conditional preferences

Michael Bernreiter*, Wolfgang Dvořák and Stefan Woltran

Institute of Logic and Computation, TU Wien, Austria

E-mail: mbernrei@dbai.tuwien.ac.at

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. Moreover, we show that in some cases the introduction of conditional preferences leads to an increase in computational complexity.

Keywords: Abstract argumentation, conditional preferences, principles, complexity

1. Introduction

Preferences are of great importance in computer science and artificial intelligence [16,25,39], where research areas such as recommender systems [18], computational social choice [31], and non-monotonic reasoning [15,40] are concerned with the representation and processing of preferences. Moreover, many situations require the use of *conditional preferences*, where a choice between two options (e.g. whether to drink tea or coffee) is dependent on other factors (e.g. the time of day). This has led to the introduction of formalisms such as CP-nets [19], which are explicitly defined to deal with conditional preferences.

In formal argumentation theory, preferences have been studied from various points of view, be it in terms of argument strength [3,4,6,32,33], preferences between values [8,12], or weighted arguments/attacks [17]. 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 [27]. There, argumentation frameworks (AFs) are built from defeasible knowledge bases containing preference rules of the form $a_1, \dots, a_n \rightarrow d_0 > d_1$, where d_0 and d_1 are defeasible rules. Similarly, there is only one recent paper we are aware of that deals with conditional preferences on the abstract level [2]. This is in contrast to unconditional preferences, which are extensively studied both in structured [36–38] and abstract [1,8,13,17,32] argumentation in the literature.

Outside of argumentation, conditional preferences appear in many situations and formalisms. One example for this are CP-nets [19], which use graphs for preference representation. Another instance is logic programming, where conditional preferences may occur in the head of rules [21,23,24] or as dedicated preference statements [20]. To demonstrate the importance of conditional preferences in common reasoning tasks, we now adapt an example given by Dung et al. [27]:

*Corresponding author. E-mail: mbernrei@dbai.tuwien.ac.at.

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.
- 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 .

In this paper, we aim to capture conditional preferences in argumentation on the abstract level rather than the structured 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 [32], 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 investigate the following topics relevant to CPAFs:

- We show that CPAFs generalize Preference-based AFs (PAFs), and demonstrate that they are capable of dealing with conditional preferences in a general manner.
- We conduct a principle-based analysis of CPAF-semantics and show that especially complete and stable semantics preserve properties that hold on PAFs. This analysis is helpful when aiming to understand the behavior of CPAF-semantics in a general manner, and lets us pinpoint differences to AFs/PAFs formally.
- We analyze the computational complexity of CPAFs in detail, showing that for some semantics (naive, complete, grounded, preferred) the introduction of conditional preferences can cause a rise in complexity compared to AFs. This gives insights into the expressiveness of CPAFs, and differentiates them further from AFs/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) [8,12] in a straightforward way, and compare CPAFs to Extended AFs (EAFs) [10,28,34] in order to highlight similarities and differences. Moreover, we discuss a recently introduced alternative approach to conditional preferences in abstract argumentation [2] and compare it to our CPAFs.

The remainder of this paper is structured as follows: Section 2 covers the necessary preliminaries on abstract argumentation. In Section 3 we introduce CPAFs and investigate them with respect to some basic properties. Section 4 contains our principle-based analysis, and in Section 5 we analyze the computational complexity of CPAFs. We discuss related formalisms in Section 6 and conclude in Section 7.

This paper is an extended version of [14]. In Section 4 we now generalize and investigate all ten principles of Kaci et al. [32] instead of just the initial six [33]. The complexity analysis featured in Section 5 is entirely new. In Section 6.1 we now provide a second translation from VAFs to CPAFs. The comparison of our formalism with lifting-based CPAFs in Section 6.3 is new as well. Moreover, this version contains additional figures and explanations to improve readability.

2. Preliminaries

We first define (abstract) argumentation frameworks [26].

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 [9]. We consider for σ the functions *cf* (conflict-free), *naive*, *adm* (admissible), *com* (complete), *grd* (grounded), *prf* (preferred), and *stb* (stable).

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

- $S \in naive(F)$ iff there is no $T \in cf(F)$ with $S \subset T$;
- $S \in adm(F)$ iff each $a \in S$ is defended by S in F ;
- $S \in com(F)$ iff $S \in adm(F)$ and each $a \in A$ defended by S in F is contained in S ;
- $S \in grd(F)$ iff $S \in com(F)$ and there is no $T \in com(F)$ with $T \subset S$;
- $S \in prf(F)$ iff $S \in adm(F)$ and there is no $T \in adm(F)$ with $S \subset T$;
- $S \in stb(F)$ iff each $a \in A \setminus S$ is attacked by S in F .

The computational complexity of AFs has been extensively studied in the literature [29], with the three central problems being those of credulous acceptance, skeptical acceptance, and verification.

Definition 3. Given an AF-semantics σ we define the following decision problems:

- *Credulous Acceptance* ($Cred_\sigma^{AF}$): given an AF F and an argument x , is $x \in S$ for some $S \in \sigma(F)$?
- *Skeptical Acceptance* ($Skept_\sigma^{AF}$): given an AF F and an argument x , is $x \in S$ in all $S \in \sigma(F)$?
- *Verification* (Ver_σ^{AF}): given an AF F and a set of arguments S , is $S \in \sigma(F)$?

Table 1 shows the complexity of these problems [29]. We assume familiarity with the complexity classes of P, NP, and coNP. Furthermore, the class Σ_2^P contains exactly those problems that can be solved in NP-time with access to an NP-oracle and Π_2^P contains the complementary problems of Σ_2^P [7].

Preference-based AFs enrich standard AFs with preferences between arguments [3,4,6,32,33].

Definition 4. A preference-based AF (PAF) is a triple $F = (A, R, \succ)$ where (A, R) is an AF and \succ is an 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

Table 1
Complexity of AFs [29]

| σ | $Cred_{\sigma}^{AF}$ | $Skept_{\sigma}^{AF}$ | Ver_{σ}^{AF} |
|--------------|----------------------|-----------------------|---------------------|
| <i>cf</i> | in P | trivial | in P |
| <i>naive</i> | in P | in P | in P |
| <i>adm</i> | NP-c | trivial | in P |
| <i>com</i> | NP-c | P-c | in P |
| <i>grd</i> | P-c | P-c | P-c |
| <i>stb</i> | NP-c | coNP-c | in P |
| <i>prf</i> | NP-c | Π_2^P -c | coNP-c |




Fig. 1. PAF F and its preference reduces from Example 2.

in the literature [32]. 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 deal with these critical attacks, e.g., by removing or reverting them.

Definition 5. Given a PAF $F = (A, R, \succ)$, a corresponding AF $\mathcal{R}_i(F) = (A, R')$ is constructed via Reduction i , where $i \in \{1, 2, 3, 4\}$, as follows:

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

For each AF-semantic σ we define the PAF-semantic σ_p^i as follows: $\sigma_p^i(F) = \sigma(\mathcal{R}_i(F))$.

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: if the weaker argument attacks the stronger argument, but not vice versa, then we retain the critical attack (as in Reduction 3) but also add a reverse attack (as in Reduction 2). 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 $\mathcal{R}_i(F)$, $i \in \{1, 2, 3, 4\}$. It can be checked that, for Reduction 1, $adm_p^1(F) = adm(\mathcal{R}_1(F)) = \{\emptyset, \{b\}, \{c\}, \{b, c\}\}$ 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 $com_p^2(F) = prf_p^2(F) = stb_p^2(F) = \{\{b\}\}$.

PAFs have the same complexity as AFs with respect to the decision problems of Definition 3: hardness results follows from the fact that PAFs generalize AFs, and membership results from the fact that the four preference reductions can be carried out in polynomial time.

A principle-based analysis of the four preference reductions was conducted for complete, grounded, preferred, and stable semantics [32,33]. To this end, ten PAF-properties were laid out and investigated. We now recall them in Definitions 6, 7, 9, 11 according to [32].

Definition 6. Let σ_p^i be a PAF-semantics. Let $\succ, \succ' \subseteq (A \times A)$ such that $\succ \cup \succ'$ is asymmetric.

- *P1* (conflict-freeness): If $(x, y) \in R$ there is no $S \in \sigma_p^i(A, R, \succ)$ such that $\{x, y\} \subseteq S$.
- *P2* (preference selects extensions 1): $\sigma_p^i(A, R, \succ \cup \succ') \subseteq \sigma_p^i(A, R, \succ)$.
- *P3* (preference selects extensions 2): $\sigma_p^i(A, R, \succ) \subseteq \sigma_p^i(A, R, \emptyset)$.

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* states that this is also true if we add preferences to a framework without any preferences, i.e., *P3* is a special case of *P2*.

Definition 7. Let σ_p^i be a PAF-semantics. Let $\succ, \succ' \subseteq (A \times A)$ such that $\succ \cup \succ'$ is asymmetric.

- *P4* (extension refinement): for all $S' \in \sigma_p^i(A, R, \succ \cup \succ')$ there is $S \in \sigma_p^i(A, R, \succ)$ such that $S \subseteq S'$.
- *P5* (extension growth): $\bigcap (\sigma_p^i(A, R, \succ)) \subseteq \bigcap (\sigma_p^i(A, R, \succ \cup \succ'))$.
- *P6* (number of extensions): $|\sigma_p^i(A, R, \succ \cup \succ')| \leq |\sigma_p^i(A, R, \succ)|$.

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.

For the next two principles, we need to define the notion of an argument's status.

Definition 8. Let $F = (A, R, \succ)$ be a PAF and $x \in A$. We write

- $status(x, F) = sk-cr$ iff x is skeptically and credulously accepted in F ;
- $status(x, F) = cr$ iff x is credulously but not skeptically accepted in F ;
- $status(x, F) = rej$ iff $status(x, F) \notin \{sk-cr, cr\}$.

We define the order over these statuses as follows: $sk-cr > cr > rej$.

Note that in stable semantics an argument is not always credulously accepted if it is skeptically accepted, since there are AFs without stable extensions. Thus, some argument x might be skeptically accepted with respect to stable semantics, yet we still might have $status(x, F) = rej$.

Definition 9. Let σ_p^i be a PAF-semantics.

- *P7* (status conservation): $status(x, (A, R, \succ \cup \{(x, y)\})) \geq status(x, (A, R, \succ))$.
- *P8* (preference-based immunity): if $(x, x) \notin R$ and $x \succ y$ for all $y \in A \setminus \{x\}$ then $status(x, (A, R, \succ)) \neq rej$.

If a semantics satisfies *P7* then the status of an argument x can not be lowered by adding a preference $x \succ y$ where x is the preferred (stronger) argument. *P8* states that if an argument x is not self-attacking and also stronger than all other arguments, then x can not be rejected.

For principles *P9* and *P10* we need the concept of paths between two arguments, by which we mean a path in the underlying undirected graph of a PAF.

Table 2

Satisfaction of various PAF-principles [32,33]. 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

| | \mathcal{R}_1 | \mathcal{R}_2 | \mathcal{R}_3 | \mathcal{R}_4 |
|--|-----------------|-----------------|-----------------|-----------------|
| $P1$ (conflict-freeness) | \times | $CGPS$ | $CGPS$ | $CGPS$ |
| $P2$ (preference selects extensions) | \times | \times | CS | \times |
| $P3$ (preference selects extensions 2) | \times | \times | CS | \times |
| $P4$ (extension refinement) | \times | \times | CGS | \times |
| $P5$ (extension growth) | \times | \times | CG | \times |
| $P6$ (number of extensions) | G | G | $CGPS$ | G |
| $P7$ (status conservation) | $CGPS$ | $CGPS$ | $CGPS$ | $CGPS$ |
| $P8$ (preference-based immunity) | CGP | CGP | \times | CPS |
| $P9$ (path preference influence 1) | $CGPS$ | $CGPS$ | $CGPS$ | $CGPS$ |
| $P10$ (path preference influence 2) | $CGPS$ | $CGPS$ | $CGPS$ | $CGPS$ |

Definition 10. Let $F = (A, R, \succ)$ be a PAF. Let $R^- = \{(x, y) | (y, x) \in R\}$. There is a path between $x \in A$ and $y \in A$ iff there is a sequence of arguments $z_1, \dots, z_n \in A$ such that $z_1 = x$, $z_n = y$, and $(z_k, z_{k+1}) \in R \cup R^-$ for all $1 \leq k < n$.

Definition 11. Let σ_p^i be a PAF-semantics.

- $P9$ (path preference influence 1): if there is no path from $x \in A$ to $y \in A$ in (A, R, \succ) then $\sigma_p^i(A, R, \succ) = \sigma_p^i(A, R, \succ \cup \{(x, y)\})$.
- $P10$ (path preference influence 2): if $(x, y) \notin R$ and $(y, x) \notin R$ then $\sigma_p^i(A, R, \succ) = \sigma_p^i(A, R, \succ \cup \{(x, y)\})$.

If $P9$ is satisfied then adding a preference between two arguments x and y that do not occur in the same weakly connected component does not change the extensions of a PAF. $P10$ is similar to $P9$, but only requires that there is no direct connection between x and y .

Table 2 shows which semantics satisfy which principle, as investigated in [32,33].

3. Conditional preference-based argumentation frameworks

As argued in the introduction, our aim is to provide a framework for reasoning with 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, 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 12. A Conditional PAF (CPAF) is a triple $F = (A, R, cond)$, where (A, R) is an AF and $cond: 2^A \rightarrow 2^{(A \times A)}$ is a function that maps each set of arguments $S \subseteq A$ to an irreflexive and asymmetric binary relation \succ_S over A .

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 $\varphi \Rightarrow x \succ y$ where φ is a propositional formula over the arguments. Indeed, this representation will be used by us in Section 5 where we analyze the complexity of CPAFs.

Just as in PAFs, preferences in CPAFs are resolved with the help of the four preference-reductions (cf. Definition 5). A set of arguments S is an extension of some CPAF if it is an extension relative to its associated preference relation $\text{cond}(S)$.

Definition 13. Let $F = (A, R, \text{cond})$ 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 $\mathcal{R}_i^S(F) = \mathcal{R}_i(A, R, \text{cond}(S))$. Given an AF-semantics σ we define the CPAF-semantics σ_{cp}^i as follows: $S \in \sigma_{cp}^i(F)$ iff $S \in \sigma(\mathcal{R}_i^S(F))$.

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 $F = (\{c_1, c_2, m_1, a_1\}, \{(c_1, c_2), (c_2, c_1)\}, \text{cond})$ with cond such that $c_1 \succ_S c_2$ iff $m_1 \in S$ but $a_1 \notin S$, $c_2 \succ_S c_1$ iff $a_1 \in S$, and $\text{cond}(S) = \emptyset$ for all other $S \subseteq A$. Figure 2 depicts F and Fig. 3 shows 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 $\sigma \in \{\text{grd}, \text{com}, \text{prf}, \text{stb}\}$ and we can conclude that $\sigma_{cp}^i(F) = \{\{m_1, a_1, c_2\}\}$.

Note that, according to Definition 13, preferred semantics 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 the S -reduct of F (and therefore $S \notin \text{prf}_{cp}^i(F)$). But this T does not have to be admissible in F , since it might not be admissible in the T -reduct of F . The situation is analogous for naive semantics. The following alternative semantics may be considered more natural:

Definition 14. Let $F = (A, R, \text{cond})$ be a CPAF and let $S \subseteq A$. Then

- $S \in \text{naive-glb}_{cp}^i(F)$ iff $S \in \text{cf}_{cp}^i(F)$ and there is no T such that $S \subset T$ and $T \in \text{cf}_{cp}^i(F)$;
- $S \in \text{prf-glb}_{cp}^i(F)$ iff $S \in \text{adm}_{cp}^i(F)$ and there is no T such that $S \subset T$ and $T \in \text{adm}_{cp}^i(F)$.

Intuitively, naive-glb_{cp}^i and prf-glb_{cp}^i maximize *globally* over all admissible sets of a CPAF, while naive_{cp}^i and prf_{cp}^i maximize *locally* over the admissible sets of the given S -reduct.




Fig. 2. The CPAF F from Example 3.




Fig. 3. The preference-reducts of the CPAF F from Fig. 2/Example 3.

Example 4. Let F be the CPAF from Example 3 and recall that $\text{prf}_{cp}^i(F) = \{\{m_1, a_1, c_2\}\}$. 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, $\text{prf-glb}_{cp}^i(F) = \{\{m_1, a_1, c_2\}, \{m_1, c_1\}\}$.

The difference between local and global maximization is not only philosophical, but impacts fundamental properties for maximization-based semantics such as I-maximality [11]. A semantics σ_{cp}^i is I-maximal if and only if $S \subseteq T$ implies $S = T$ for all CPAFs F and all $S, T \in \sigma_{cp}^i(F)$.

Proposition 1. prf-glb_{cp}^i is I-maximal, but prf_{cp}^i is not, where $i \in \{1, 2, 3, 4\}$.

Proof. I-maximality of prf-glb_{cp}^i follows from Definition 14. Regarding counter examples for prf_{cp}^i we consider the preference-reductions separately. Reduction 1: consider the CPAF depicted in Fig. 4a, i.e., $F = (\{a, b\}, \{(a, b)\}, \text{cond})$ with cond such that $b \succ_{\{a,b\}} a$. Then $\{a\} \in \text{prf}_{cp}^1(F)$ and $\{a, b\} \in \text{prf}_{cp}^1(F)$. Reductions 2 and 4: consider the CPAF depicted in Fig. 4b, i.e., $F' = (\{a, b, c\}, \{(a, b), (b, c), (c, a)\}, \text{cond})$ with cond such that $a \succ_{\{a\}} c$. Then $\emptyset \in \text{prf}_{cp}^i(F')$ and $\{a\} \in \text{prf}_{cp}^i(F')$. Reduction 3: consider the CPAF depicted in Fig. 4c, i.e., $F'' = (\{a, b, c\}, \{(a, b), (b, a), (b, c), (c, a)\}, \text{cond})$ with cond such that $a \succ_{\emptyset} b$. Then $\emptyset \in \text{prf}_{cp}^3(F'')$ and $\{b\} \in \text{prf}_{cp}^3(F'')$. \square

One may be tempted to deduce from the above proposition that prf-glb_{cp}^i is more suitable as a default preferred semantics than prf_{cp}^i . However, we will see in Section 6.1 that prf_{cp}^i 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 local and global subset maximization. Like preferred semantics, naive and stable semantics satisfy I-maximality on AFs. Interestingly, on CPAFs, this depends on the preference-reduction.

Proposition 2. naive-glb_{cp}^i is I-maximal for $i \in \{1, 2, 3, 4\}$. Moreover, naive_{cp}^j is I-maximal for $j \in \{2, 3, 4\}$ but not for $j = 1$.

Proof. I-maximality of naive-glb_{cp}^i follows from Definition 14. I-maximality of naive_{cp}^j with $j \in \{2, 3, 4\}$ follows from the fact that Reductions 2, 3, and 4 do not remove conflicts between arguments, and therefore conflict-free sets are the same across all S -reducts. For naive_{cp}^1 we can use the same counterexample as for prf_{cp}^1 (cf. Proposition 1 and Fig. 4a). \square

Proposition 3. stb_{cp}^j is I-maximal for $j \in \{2, 3, 4\}$ but not for $j = 1$.

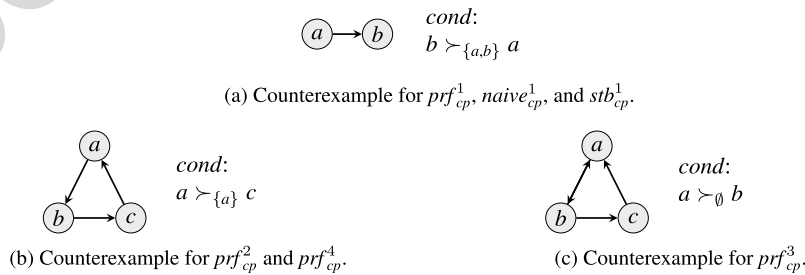


Fig. 4. Counterexamples for I-maximality (cf. Propositions 1,2,3).

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

A further well-known property of AFs is that if an argument set S is stable in a framework F , then S is also preferred in F [26]. The same is true for CPAFs, with the exception of preferred semantics utilizing global maximization and Reduction 1.

Proposition 4. *If $S \in stb_{cp}^i(F)$ then $S \in prf_{cp}^i(F)$ for $i \in \{1, 2, 3, 4\}$. Moreover, if $S \in stb_{cp}^j(F)$ then $S \in prf-glb_{cp}^j(F)$ for $j \in \{2, 3, 4\}$. However, $S \in stb_{cp}^1(F)$ does not necessarily imply $S \in prf-glb_{cp}^1(F)$.*

Proof. Let $F = (A, R, cond)$ be a CPAF, and let $S \in stb_{cp}^i(F)$, where $i \in \{1, 2, 3, 4\}$. Then $S \in stb(\mathcal{R}_i^S(F))$. Since $\mathcal{R}_i^S(F)$ is an AF this implies $S \in prf(\mathcal{R}_i^S(F))$ which means that $S \in prf_{cp}^i(F)$.

Now let $j \in \{2, 3, 4\}$. If $S \in stb_{cp}^j(F)$ then every argument in $\mathcal{R}_j^S(F)$ is either in S or attacked by it. Towards a contradiction, assume there is $T \in adm_{cp}^j(F)$ such that $T \supset S$. Then there is some $x \in T \setminus S$. Since S attacks x in $\mathcal{R}_j^S(F)$ there is a conflict between some $y \in S$ and x in the underlying AF (A, R) of F . Note that $y \in T$. But Reductions 2, 3, 4 can not remove conflicts between arguments, i.e., $T \notin cf(\mathcal{R}_j^T(F))$. Contradiction.

For $prf-glb_{cp}^1$, let F be the CPAF from Fig. 4a). Then $\{a\} \in stb_{cp}^1(F)$ but $\{a\} \notin prf-glb_{cp}^1(F)$. \square

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 $\{a\} \in grd_{cp}^2(F)$ and $\{b\} \in grd_{cp}^2(F)$. 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 every CPAF-semantics considered here generalizes their corresponding PAF-semantics, i.e., that CPAFs generalize PAFs.

Proposition 5. *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 $\sigma \in \{cf, naive, adm, com, grd, prf, stb\}$. Then $\sigma_{cp}^i(F) = \sigma_p^i(A, R, \succ)$. Furthermore, $naive-glb_{cp}^i(F) = naive_p^i(A, R, \succ)$ and $prf-glb_{cp}^i(F) = prf_p^i(A, R, \succ)$.*

4. Principle-based analysis

Principles play an important role in argumentation theory, as they allow us to examine the vast amount of semantics defined for AFs in a general way [11,30,41]. In this section, we generalize the principles of Kaci et al. [32] for PAFs (cf. Definitions 6, 7, 9, 11) to account for conditional preferences. We then investigate by which semantics these principles are satisfied, and show that there are differences to PAFs.

In the case of PAFs, adding more preferences to a framework (A, R, \succ) means that we now deal with the PAF $(A, R, \succ \cup \succ')$. In the case of CPAFs, if we want $(A, R, cond')$ to have at least the same preferences as $(A, R, cond)$, we must require that $cond(S) \subseteq cond'(S)$ for all $S \subseteq A$. However, if we only

want to add a single preference $x \succ y$ to a CPAF, then we add $x \succ_S y$ to only a single subset S and leave the preferences associated with other subsets unchanged. Given the above considerations, generalizing the PAF-principles to CPAF-principles is quite straightforward. The notions of an argument's status and paths between two arguments in a CPAF are defined analogously to PAFs (cf. Definitions 8, 10), e.g., $status(x, F) = cr$ iff x is contained in some but not all extensions of the CPAF F .

Definition 15. Let σ_{cp}^i be a CPAF-semantic. 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 $S \subseteq A$. Moreover, $cond_{(x,y)}$ is the same as $cond$ except that for some $S \subseteq A$ we have $(x, y) \in cond_{(x,y)}(S)$ but $(x, y), (y, x) \notin cond(S)$. Lastly, $cond_{\emptyset}(S) = \emptyset$ for all $S \subseteq A$.

- $P1^*$ (conflict-freeness): If $(x, y) \in R$ there is no $S \in \sigma_{cp}^i(A, R, cond)$ such that $\{x, y\} \subseteq S$.
- $P2^*$ (preference selects extensions): $\sigma_{cp}^i(A, R, cond') \subseteq \sigma_{cp}^i(A, R, cond)$.
- $P3^*$ (preference selects extensions 2): $\sigma_{cp}^i(A, R, cond) \subseteq \sigma_{cp}^i(A, R, cond_{\emptyset})$.
- $P4^*$ (extension refinement): for all $S' \in \sigma_{cp}^i(A, R, cond')$ there is $S \in \sigma_{cp}^i(A, R, cond)$ s.t. $S \subseteq S'$.
- $P5^*$ (extension growth): $\bigcap(\sigma_{cp}^i(A, R, cond)) \subseteq \bigcap(\sigma_{cp}^i(A, R, cond'))$.
- $P6^*$ (number of extensions): $|\sigma_{cp}^i(A, R, cond')| \leq |\sigma_{cp}^i(A, R, cond)|$.
- $P7^*$ (status conservation): $status(x, (A, R, cond_{(x,y)})) \geq status(x, (A, R, cond))$.
- $P8^*$ (preference-based immunity): if $(x, x) \notin R$ and if $cond$ is defined such that for all $S \subseteq A$ and all $y \in A \setminus \{x\}$ we have $x \succ_S y$ then $status(x, (A, R, cond)) \neq rej$.
- $P9^*$ (path preference influence 1): if there is no path from $x \in A$ to $y \in A$ in $(A, R, cond)$ then $\sigma_{cp}^i(A, R, cond) = \sigma_{cp}^i(A, R, cond_{(x,y)})$.
- $P10^*$ (path preference influence 2): if $(x, y) \notin R$ and $(y, x) \notin R$ then $\sigma_{cp}^i(A, R, cond) = \sigma_{cp}^i(A, R, cond_{(x,y)})$.

The following lemma establishes some relationships between the CPAF-principles and is a generalization of known relationships between PAF-principles [32].

Lemma 6. *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 $P5^*$.*

Proof. For $P3^*$, $P4^*$, and $P6^*$ this is easy to see. For $P5^*$ we argue this in detail: by contrapositive, let σ_{cp}^i be a semantic that always returns at least one extension, but does not satisfy $P5^*$. Thus, there must be $A, R, cond, cond'$ with $cond(S) \subseteq cond'(S)$ for all $S \subseteq A$, such that $\bigcap(\sigma_{cp}^i(A, R, cond)) \not\subseteq \bigcap(\sigma_{cp}^i(A, R, cond'))$. Then there is $x \in A$ such that $x \in \bigcap(\sigma_{cp}^i(A, R, cond))$ but $x \notin \bigcap(\sigma_{cp}^i(A, R, cond'))$, i.e., there is $E \subseteq A$ such that $x \notin E$ and $E \in \sigma_{cp}^i(A, R, cond')$. Of course, $E \notin \sigma_{cp}^i(A, R, cond)$, otherwise $x \notin \bigcap(\sigma_{cp}^i(A, R, cond))$. But then $\sigma_{cp}^i(A, R, cond') \not\subseteq \sigma_{cp}^i(A, R, cond)$, i.e., σ_{cp}^i does not satisfy $P2^*$. \square

Observe that, since CPAFs are a generalization of PAFs (cf. Proposition 5), a CPAF-semantic σ_{cp}^i can not satisfy Pj^* if the corresponding PAF-semantic 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 7. *If σ_p^i satisfies $P2$ then σ_{cp}^i satisfies $P2^*$.*

Proof. By contrapositive, 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 $E \in \sigma(\mathcal{R}_i(A, R, cond'(E)))$ but $E \notin \sigma(\mathcal{R}_i(A, R, cond(E)))$, i.e., σ_p^i does not satisfy $P2$. \square

Lemma 7 implies that complete and stable semantics satisfy $P2^*$ on CPAFs under Reduction 3. By Lemma 6 these semantics also satisfy $P3^*$, $P4^*$, and $P6^*$. However, we can not use Lemma 6 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. Counterexamples for the satisfaction of various principles can also be found for grounded semantics, both variants of the preferred semantics, and even stable semantics in the case of $P8^*$.

Lemma 8. *The following holds:*

- grd_{cp}^i , where $i \in \{1, 2, 3, 4\}$, does not satisfy any of $P4^*$, $P5^*$, or $P6^*$;
- com_{cp}^3 does not satisfy $P5^*$;
- prf_{cp}^3 and $prf-glb_{cp}^3$ do not satisfy $P6^*$;
- σ_{cp}^i , where for $\sigma \in \{com, grd, prf, stb\}$ and $i \in \{1, 2, 3, 4\}$, does not satisfy $P8^*$.

Proof. We provide counterexamples for all cases.

- For grd_{cp}^i and $P4^*$, $P5^*$, $P6^*$, consider $A = \{a, b\}$, $R = \{(a, b), (b, a)\}$, and $cond/cond'$ as shown in Fig. 5a. Then $grd_{cp}^i(A, R, cond) = \{\{a\}\}$ while $grd_{cp}^i(A, R, cond') = \{\{a\}, \{b\}\}$.
- For com_{cp}^3 and $P5^*$, consider $A = \{a, b\}$, $R = \{(a, b), (b, a)\}$, and $cond/cond'$ as shown in Fig. 5b. Then $com_{cp}^3(A, R, cond) = \{\{a\}\}$ while $com_{cp}^3(A, R, cond') = \emptyset$.
- For prf_{cp}^3 , $prf-glb_{cp}^3$ and $P6^*$, consider $A = \{a, b, c, d\}$, $R = \{(a, c), (c, a), (b, d), (d, b), (c, c), (d, d)\}$, and $cond/cond'$ as shown in Fig. 5c. Then $prf_{cp}^3(A, R, cond) = prf-glb_{cp}^3(A, R, cond) = \{\{a, b\}\}$ while $prf_{cp}^3(A, R, cond') = prf-glb_{cp}^3(A, R, cond') = \{\{a\}, \{b\}\}$.
- Regarding $P8^*$, consider the CPAF $F = (A, R, cond)$ shown in Fig. 5d. Note that $(a, a) \notin R$ and $a \succ_S y$ for all $S \subseteq A$ and all $y \notin \{a\}$. Observe that b is unattacked in $\mathcal{R}_i^{(a)}(F)$. Thus,

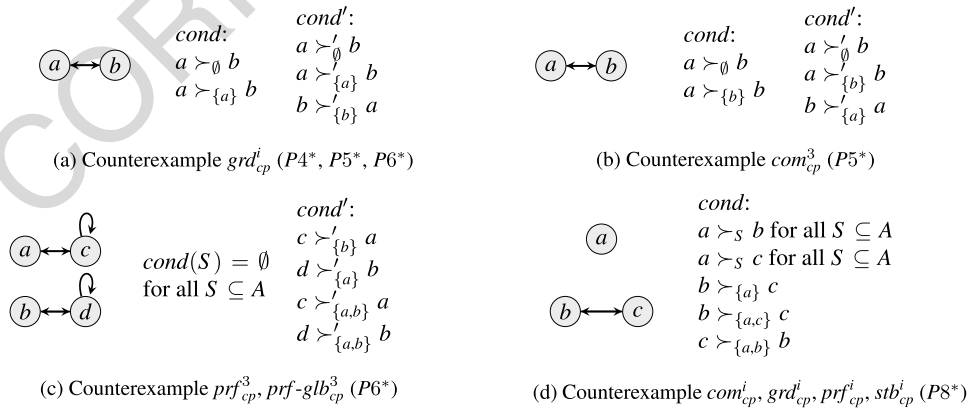


Fig. 5. Counterexamples used in Lemma 8.

$\{a\} \notin \sigma(\mathcal{R}_i^{(a)}(F))$ for $\sigma \in \{com, grd, prf, stb\}$. Moreover, b is not defended against c in $\mathcal{R}_i^{(a,b)}(F)$. Analogously for c in $\mathcal{R}_i^{(a,c)}(F)$. Thus, $status(a, F) = rej$. \square

We have now fully investigated the first six principles. It remains to examine principles 7-10, of which we so far only know that $P8^*$ is not satisfied by most semantics. It turns out that $P8^*$ is retained when using preferred semantics with global maximization.

Lemma 9. $prf\text{-}glb_{cp}^i$ satisfies $P8^*$ for $i \in \{1, 2, 4\}$.

Proof. Let $F = (A, R, cond)$ be a CPAF containing an argument $x \in A$ such that $(x, x) \notin R$ and for all $S \subseteq A$ and all $y \in A \setminus \{x\}$ we have $x \succ_S y$. Specifically, this means that $x \succ_{\{x\}} y$ for all $y \in A \setminus \{x\}$. Then, by definition of Reduction $i \in \{1, 2, 4\}$, x defends itself against all attacks in $\mathcal{R}_i^{(x)}(F)$. Thus, $\{x\} \in adm_{cp}^i(F)$. By this and the definition of $prf\text{-}glb$, there is some $E \in prf\text{-}glb_{cp}^i(F)$ such that $x \in E$. \square

Now we turn our attention to $P7^*$, where, analogously to $P2^*$ (cf. Lemma 7), it turns out that satisfaction carries over from PAFs to CPAFs.

Lemma 10. If σ_p^i satisfies $P7$ then σ_{cp}^i satisfies $P7^*$.

Proof. By contrapositive, assume σ_{cp}^i does not satisfy $P7^*$. Then there is a CPAF $F = (A, R, cond)$ such that $status(x, (A, R, cond_{(x,y)})) < status(x, (A, R, cond))$. This means there is some $S \subseteq A \cup \{x\}$ for which $S \in \sigma_{cp}^i(A, R, cond)$ but $S \notin \sigma_{cp}^i(A, R, cond_{(x,y)})$. By the definition of CPAF-semantics this means that $S \in \sigma_p^i(A, R, cond(S))$ but $S \notin \sigma_p^i(A, R, cond(S) \cup \{(x, y)\})$, i.e., σ_p^i does not satisfy $P7$. \square

However, unlike in the case of $P2^*$, the above lemma does not constitute an exhaustive investigation of $P7^*$. The reason for this is that $P7$, in contrast to $P2$, is satisfied by preferred semantics on PAFs (cf. Table 2). Lemma 10 only allows us to conclude that prf_{cp}^i satisfies $P7^*$, but it says nothing about the satisfaction of $prf\text{-}glb_{cp}^i$. We find that $P7^*$ is satisfied also when maximizing admissible sets globally.

Lemma 11. $prf\text{-}glb_{cp}^i$ satisfies $P7^*$.

Proof. Let $F = (A, R, cond)$ be an arbitrary CPAF and $x \in A$. Let $F_{(x,y)} = (A, R, cond_{(x,y)})$ as specified in Definition 15. There are three possible cases:

- (1) $status(x, F) = rej$. Then $status(x, F_{(x,y)}) \geq status(x, F)$ trivially holds.
- (2) $status(x, F) = cr$. Then there is some $S \in adm_{cp}^i(F)$ with $x \in S$. We distinguish two cases:
 - (a) $cond(S) = cond_{(x,y)}(S)$. Then clearly $S \in adm_{cp}^i(F_{(x,y)})$.
 - (b) $cond(S) \neq cond_{(x,y)}(S)$. Then $cond_{(x,y)}$ is the same as $cond$ except that $(x, y) \in cond_{(x,y)}(S)$ but $(x, y), (y, x) \notin cond(S)$ for some $y \in A \setminus \{x\}$. Adding the preference $x \succ y$ via $cond_{(x,y)}$ does not introduce any new attacks against S in $\mathcal{R}_i^S(F_{(x,y)})$, no matter which of the preference reductions we consider. Thus, $S \in adm_{cp}^i(F_{(x,y)})$.

In both cases we have $S \in adm_{cp}^i(F_{(x,y)})$ and therefore $T \in prf\text{-}glb_{cp}^i(F_{(x,y)})$ for some $T \supseteq S$. Thus, $status(x, F_{(x,y)}) \geq cr$.

Table 3

Satisfaction of CPAF-principles. C stands for complete, G for grounded, P for preferred (both prf_{cp}^i and $prf-glb_{cp}^i$), and S for stable. P_g indicates that $prf-glb_{cp}^i$ satisfies the principle but prf_{cp}^i does not. If a cell is marked with \times then none of the investigated semantics satisfy this principle

| | \mathcal{R}_1 | \mathcal{R}_2 | \mathcal{R}_3 | \mathcal{R}_4 |
|--|-----------------|-----------------|-----------------|-----------------|
| $P1^*$ (conflict-freeness) | \times | $CGPS$ | $CGPS$ | $CGPS$ |
| $P2^*$ (preference selects extensions) | \times | \times | CS | \times |
| $P3^*$ (preference selects extensions 2) | \times | \times | CS | \times |
| $P4^*$ (extension refinement) | \times | \times | CS | \times |
| $P5^*$ (extension growth) | \times | \times | \times | \times |
| $P6^*$ (number of extensions) | \times | \times | CS | \times |
| $P7^*$ (status conservation) | $CGPS$ | $CGPS$ | $CGPS$ | $CGPS$ |
| $P8^*$ (preference-based immunity) | P_g | P_g | \times | P_g |
| $P9^*$ (path preference influence 1) | $CGPS$ | $CGPS$ | $CGPS$ | $CGPS$ |
| $P10^*$ (path preference influence 2) | $CGPS$ | $CGPS$ | $CGPS$ | $CGPS$ |

(3) $status(x, F) = sk-cr$. By the same line of reasoning as in case (2) we have $status(x, F_{(x,y)}) \geq cr$. Towards a contradiction, assume that $status(x, F_{(x,y)}) \neq sk-cr$. Then there is some $S \in prf-glb_{cp}^i(F_{(x,y)})$ such that $x \notin S$. Since $status(x, F) = sk-cr$ we know that $S \notin prf-glb_{cp}^i(F)$. One of the following must be the case:

- (a) $S \notin cf_{cp}^i(F)$. Since $S \in prf-glb_{cp}^i(F_{(x,y)})$ it must be that $S \in cf_{cp}^i(F_{(x,y)})$. Then it must be that the additional preference in $cond_{(x,y)}$ removes a conflict between two arguments in S . But this preference is $x \succ y$ for some $y \in A \setminus \{x\}$. Thus, $x \in S$. Contradiction.
- (b) $S \in cf_{cp}^i(F)$ but $S \notin adm_{cp}^i(F)$. Since $S \in prf-glb_{cp}^i(F_{(x,y)})$ it must be that $S \in adm_{cp}^i(F_{(x,y)})$. However, the additional preference $x \succ y$ added via $cond_{(x,y)}$ at most adds an attack (x, y) . Since $x \notin S$ this means that S is still not defended in $\mathcal{R}_i^S(F_{(x,y)})$. Contradiction.
- (c) $S \in adm_{cp}^i(F)$ but there is $T \supset S$ such that $T \in prf-glb_{cp}^i(F)$. Since $status(x, F) = sk-cr$ we know that $x \in T$. Since $T \in prf-glb_{cp}^i(F)$ we have $T \in adm_{cp}^i(F)$. By the same line of reasoning as in case (2) we can conclude that $T \in adm_{cp}^i(F_{(x,y)})$ and therefore $S \notin prf-glb_{cp}^i(F_{(x,y)})$. Contradiction.

In all three cases we arrive at a contradiction. Thus, $status(x, F_{(x,y)}) = sk-cr$. \square

Lastly, we must consider principles 9 and 10. Below, we show that they retain the satisfaction of all principles under all considered semantics.

Lemma 12. σ_{cp}^i satisfies $P9^*$ and $P10^*$ for $\sigma \in \{com, grd, prf, prf-glb, stb\}$.

Proof. Let $F = (A, R, cond)$ be an arbitrary CPAF and $x \in A$. Let $F_{(x,y)} = (A, R, cond_{(x,y)})$ as specified in Definition 15. Note that the premise for $P9^*$ (there is no path from $x \in A$ to $y \in A$) implies the premise of $P10^*$ ($(x, y) \notin R$ and $(y, x) \notin R$). If $(x, y) \notin R$ and $(y, x) \notin R$ then the additional preference $x \succ y$ added via $cond_{(x,y)}$ does not delete or add any attacks, regardless of which preference reduction we consider. This means that $\mathcal{R}_i^S(F) = \mathcal{R}_i^S(F_{(x,y)})$ for all $S \subseteq A$ and all $i \in \{1, 2, 3, 4\}$ and therefore $\sigma_{cp}^i(A, R, cond) = \sigma_{cp}^i(A, R, cond_{(x,y)})$ for all $i \in \{1, 2, 3, 4\}$. \square

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

Theorem 13. *The satisfaction of CPAF-principles depicted in Table 3 holds.*

To summarize, complete and stable semantics preserve the satisfaction of PAF-principles in most cases. Grounded semantics no longer satisfies any of the principles 1-6 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 8). Unlike on PAFs, complete semantics does not satisfy $P5^*$ (extension growth) under Reduction 3. Furthermore, neither variant of preferred semantics satisfies $P6^*$ (number of extensions) under Reduction 3. As for principles 7-10, we note that only $P8^*$ is no longer satisfied by all semantics.

5. Computational complexity

The computational complexity of Dung-style AFs and various generalizations thereof has received considerable attention in the literature [29]. Indeed, complexity results give insights into the expressiveness of specific argumentation formalisms and help to find appropriate methods for solving a given problem. As discussed in Section 2, AFs and PAFs have the same properties with regards to complexity, i.e., none of the four preference reductions result in a higher complexity when considering unconditional preferences in the setting of Dung-AFs. The situation is not as clear when dealing with conditional preferences. As we have seen in previous sections, CPAFs do not necessarily have unique grounded extensions, there are CPAFs without any complete extensions, and there is more than one way of dealing with subset maximization (recall the *naive/naive-glb* and *prf/prf-glb* semantics). In this section, we show that these differences between CPAFs and AFs/PAFs have an impact on complexity.

We define $Ver_{\sigma,i}^{CPAF}$, $Cred_{\sigma,i}^{CPAF}$, and $Skept_{\sigma,i}^{CPAF}$ analogously to Ver_{σ}^{AF} , $Cred_{\sigma}^{AF}$, and $Skept_{\sigma}^{AF}$ (cf. Definition 3), with the difference that the framework in question is now a CPAF instead of an AF and that we appeal to the σ_{cp}^i semantics of Definitions 13 and 14 rather than the AF-semantics of Definition 2:

Definition 16. Given a CPAF-semantics σ_{cp}^i we define the following decision problems:

- *Credulous Acceptance* ($Cred_{\sigma,i}^{CPAF}$): given a CPAF F and an argument x , is $x \in S$ for some $S \in \sigma_{cp}^i(F)$?
- *Skeptical Acceptance* ($Skept_{\sigma,i}^{CPAF}$): given a CPAF F and an argument x , is $x \in S$ in all $S \in \sigma_{cp}^i(F)$?
- *Verification* ($Ver_{\sigma,i}^{CPAF}$): given a CPAF F and a set of arguments S , is $S \in \sigma_{cp}^i(F)$?

In the interest of generality, we did not impose a specific method to represent conditional preferences in the previous sections. However, when analyzing the computational complexity of CPAFs, it is necessary to decide on more specific representations if tight bounds are to be found. Therefore, we will assume conditional preferences to be expressed succinctly as arbitrary propositional formulas. Note that if preferences would be stored explicitly for each possible set of arguments, the input size of our problems would always be exponentially larger than the underlying AF itself, and thus some decision problems for CPAFs would be in lower complexity classes than their counterparts for AFs.

Specifically, given the set of argument A in a framework $F = (A, R, cond)$ we allow a finite number of rules $\varphi \Rightarrow x \succ y$ where $x, y \in A$ and φ is a propositional formula built from atoms in A and the usual connectives (\neg, \wedge, \vee). As for the meaning of these rules, we define that for some $S \subseteq A$ we have $x \succ_S y$ iff there is a rule $\varphi \Rightarrow x \succ y$ such that $S \models \varphi$ and there is no rule $\varphi' \Rightarrow y \succ x$ such that $S \models \varphi'$.¹

¹A set of atoms S can be seen as an interpretation, with x set to true under S iff $x \in S$.

Table 4

Complexity of CPAFs with conditional preferences represented via finitely many rules of the form $\varphi \Rightarrow x > y$. Underlines indicate a rise in complexity compared to AFs

| σ | $Cred_{\sigma,1}^{CPAF} / Cred_{\sigma,j \in \{2,3,4\}}^{CPAF}$ | $Skept_{\sigma,1}^{CPAF} / Skept_{\sigma,j \in \{2,3,4\}}^{CPAF}$ | $Ver_{\sigma,1}^{CPAF} / Ver_{\sigma,j \in \{2,3,4\}}^{CPAF}$ |
|------------------|---|---|---|
| <i>cf</i> | in P | trivial | in P |
| <i>naive</i> | <u>NP-c/in P</u> | <u>coNP-c/in P</u> | in P |
| <i>naive-glb</i> | in P | <u>Π_2^P-c/in P</u> | <u>coNP-c/in P</u> |
| <i>adm</i> | NP-c | trivial | in P |
| <i>com</i> | NP-c | <u>coNP-c</u> | in P |
| <i>grd</i> | <u>NP-c</u> | <u>coNP-c</u> | P-c |
| <i>stb</i> | NP-c | coNP-c | in P |
| <i>prf</i> | <u>Σ_2^P-c</u> | Π_2^P -c | coNP-c |
| <i>prf-glb</i> | <u>NP-c</u> | Π_2^P -c | coNP-c |

Observe that, given $S \subseteq A$, it is possible to compute $>_S$ in polynomial time with respect to the size of the given framework F since $S \models \varphi$ can be decided in polynomial time for each rule $\varphi \Rightarrow x > y$.²

Our complexity results are summarized in Table 4. Note that problems for *naive/naive-glb* semantics become harder only under Reduction 1. Intuitively, this is because Reduction 1 can remove conflicts between arguments altogether, unlike Reductions 2-4. Observe that *naive-glb* under Reduction 1 is the only semantics for which the verification problem becomes harder (coNP-complete) compared to AFs (in P). As a result, skeptical acceptance for *naive-glb* is Π_2^P -complete, i.e., the complexity rises by two levels in the polynomial hierarchy compared to the case of AFs. For complete semantics, skeptical acceptance is now coNP-complete regardless of which preference reduction is used. With respect to grounded semantics we see an increase in complexity for both credulous acceptance (NP-complete) and skeptical acceptance (coNP-complete). Lastly, for preferred semantics with local maximization, credulous acceptance rises by one level in the polynomial hierarchy compared to AFs.

Theorem 14. *The complexity results for CPAFs depicted in Table 4 hold.*

The remainder of this section is dedicated to proving Theorem 14. We first consider the verification problem, for which most semantics have the same complexity as in the case of AFs.

Lemma 15. *$Ver_{\sigma,i}^{CPAF}$, where $i \in \{1, 2, 3, 4\}$, has the same complexity properties with regards to membership and hardness as Ver_{σ}^{AF} for $\sigma \in \{cf, naive, adm, com, grd, stb, prf\}$. Moreover, $Ver_{prf-glb,i}^{CPAF}$ is coNP-complete for $i \in \{1, 2, 3, 4\}$.*

Proof. Hardness follows from the fact that CPAFs are a generalization of AFs. Membership for $\sigma \in \{cf, naive, adm, com, grd, stb, prf\}$: given a CPAF F and a set of arguments S , we can determine $>_S$ and therefore also $\mathcal{R}_i^S(F)$ in polynomial time. It then suffices to check whether $S \in \sigma(\mathcal{R}_i^S(F))$. Membership for $\sigma = prf-glb$: let (F, S) be an arbitrary instance of $Ver_{prf-glb,i}^{CPAF}$, i.e., $F = (A, R, cond)$ is a CPAF and $S \subseteq A$ is a set of arguments. First, check in polynomial time whether $S \in adm_{cp}^i(F)$. Then, in coNP-time, check that for all T we have either $T \subseteq S$ or $T \notin adm_{cp}^i(F)$. \square

²In fact, for our membership results the explicit representation of rules using propositional formulas is not necessary. It suffices to have some representation such that, given $S \subseteq A$, we can determine $>_S$ in polynomial time with respect to the size of $F = (A, R, cond)$. However, for hardness results, a more concrete representation such as via our rules is necessary.

For naive semantics with global maximization (*naive-glb*) we see a rise in complexity, but only when using Reduction 1. The following proof makes use of Reduction 1's ability to delete conflicts between arguments. By SAT we denote the NP-complete satisfiability problem for propositional formulas, and by UNSAT we denote its complementary problem which is coNP-complete.

Lemma 16. $Ver_{naive-glb,j}^{CPAF}$ is in P for $j \in \{2, 3, 4\}$. $Ver_{naive-glb,1}^{CPAF}$ is coNP-complete.

Proof. Let (F, S) be an arbitrary instance of $Ver_{naive-glb,i}^{CPAF}$, i.e., $F = (A, R, cond)$ is a CPAF and $S \subseteq A$ is a set of arguments. For Reductions 2, 3, and 4 it suffices to check whether $S \in naive(A, R)$ since these reductions can not remove or add conflicts.

We now turn our attention to Reduction 1. coNP-membership: first check whether $S \in cf_{cp}^1(F)$. Then, in coNP-time, check that for all T we have either $T \subseteq S$ or $T \notin cf_{cp}^1(F)$. coNP-hardness by reduction from UNSAT: let φ be an arbitrary propositional formula over variables X . Let a be a fresh variable, i.e., $a \notin X$. We construct an instance $(F, \{a\})$ of $Ver_{naive-glb,1}^{CPAF}$ as follows: $F = (A, R, cond)$ with $A = X \cup \{a\}$, $R = \{(x, a) | x \in X\}$, and $cond$ defined by the rules $\varphi \Rightarrow a \succ x$ for $x \in X$, i.e., $a \succ_S x$ iff $S \models \varphi$. Figure 6 depicts the above construction. We now show that φ is unsatisfiable iff $\{a\} \in naive-glb_{cp}^1(F)$ (i.e., $(F, \{a\})$ is a yes-instance of $Ver_{naive-glb,1}^{CPAF}$).

- Suppose φ is unsatisfiable. This means that for all $x \in X$ and all $S \subseteq A$ we have $a \not\succeq_S x$, i.e., for each $x \in X$ the attack (x, a) is present in $\mathcal{R}_1^S(F)$. Thus, there is no conflict-free set containing a other than $\{a\}$ which implies $\{a\} \in naive-glb_{cp}^1(F)$.
- Suppose φ is satisfiable. Then there is an interpretation $I \subseteq X$ such that $I \models \varphi$. We assume that $I \neq \emptyset$. This is permissible since we can check in polynomial time whether \emptyset satisfies φ , and if this is the case, return a trivial no-instance of $Ver_{naive-glb,1}^{CPAF}$. Consider $S = I \cup \{a\}$. Since a does not appear in φ we have $S \models \varphi$ and therefore $a \succ_S x$ for all $x \in X$. Thus, $S \in cf_{cp}^1(F)$ and, since $\{a\} \subset S$, $\{a\} \notin naive-glb_{cp}^1(F)$. \square

We now consider credulous and skeptical acceptance, starting with semantics based solely on conflict-freeness. Let us first cover the cases in which there is no rise in complexity.

Lemma 17. $Cred_{\sigma,i}^{CPAF}$ is in P for $\sigma \in \{cf, naive-glb\}$, $i \in \{1, 2, 3, 4\}$. $Cred_{naive,j}^{CPAF}$ is in P for $j \in \{2, 3, 4\}$.

Proof. Let (F, x) be an instance of $Cred_{\sigma,i}^{CPAF}$. For $\sigma = cf$ it suffices to check whether x is self-attacking in the underlying AF of F , since self-attacks are not removed by any of the four reductions. For $\sigma = naive-glb$ it suffices to test whether (F, x) is a yes-instance of $Cred_{cf,i}^{CPAF}$. For $\sigma = naive$ and Reductions 2, 3, and 4 it is enough to check whether x appears in a naive set of the underlying AF, since these reductions can not remove conflicts. \square




Fig. 6. Construction used in the proof of Lemma 16. Given a formula φ over variables $X = \{x_1, x_2, \dots, x_n\}$, a CPAF F is constructed such that φ is unsatisfiable iff $\{a\} \in naive-glb_{cp}^1(F)$.

Lemma 18. $Skept_{cf,i}^{CPAF}$ is trivial for $i \in \{1, 2, 3, 4\}$. $Skept_{\sigma,j}^{CPAF}$ is in \mathbf{P} for $\sigma \in \{naive, naive-glb\}$ and $j \in \{2, 3, 4\}$.

Proof. Let (F, x) be an arbitrary instance of $Skept_{naive-glb,i}$, i.e., $F = (A, R, cond)$ is a CPAF and $x \in A$ is an argument. Note that \emptyset is always conflict-free in $\mathcal{R}_i^\emptyset(F)$, i.e., (F, x) is trivially a no-instance. For $\sigma \in \{naive, naive-glb\}$ and Reductions 2, 3, and 4 it is enough to solve the problem on the underlying AF, since these reductions can not remove conflicts. \square

For naive semantics with local maximization the complexity rises by one level in the polynomial hierarchy under Reduction 1.

Lemma 19. $Cred_{naive,1}^{CPAF}$ is NP-complete and $Skept_{naive,1}^{CPAF}$ is coNP-complete.

Proof. We will consider the complementary problem of $Skept_{naive,1}^{CPAF}$ and show that it is NP-complete since this allows us to prove both results simultaneously.

NP-Membership: given a CPAF $F = (A, R, cond)$ and an argument $x \in A$, guess a set $S \subseteq A$ and, in polynomial time, check whether $S \in naive_{cp}^1(F)$ and $x \in S$ (resp. $S \in naive_{cp}^1(F)$ and $x \notin S$).

NP-hardness by reduction from SAT: let φ be an arbitrary propositional formula over a set of variables X . Let a and b be fresh atoms, i.e., $a, b \notin X$. We construct an instance (F, a) of $Cred_{naive,1}^{CPAF}$ as follows: $F = (A, R, cond)$ with $A = X \cup \{\bar{x} | x \in X\} \cup \{a, b\}$, $R = \{(x, \bar{x}), (\bar{x}, x) | x \in X\} \cup \{(a, b)\}$, and $cond$ defined by the rule $\neg\varphi \wedge a \wedge \neg b \Rightarrow b \succ a$, i.e., $b \succ_S a$ iff $S \models \neg\varphi \wedge a \wedge \neg b$. The above construction is visualized in Fig. 7. We show that φ is satisfiable iff (F, a) is a yes-instance of $Cred_{naive,1}^{CPAF}$ iff (F, b) is a no-instance of $Skept_{naive,1}^{CPAF}$:

- Assume φ is satisfiable. Then there is an interpretation $I \subseteq X$ such that $I \models \varphi$. Then also $I \cup \{a\} \models \varphi$ and $I \cup \{a\} \not\models \neg\varphi \wedge a \wedge \neg b$. Let $S = I \cup \{\bar{x} | x \notin I\} \cup \{a\}$. Note that $b \not\succeq_S a$, which means that a and b are in conflict in $\mathcal{R}_1^S(F)$. Furthermore, for all $x \in X$, we have either $x \in S$ or $\bar{x} \in S$, but not both. Thus, $S \in naive(\mathcal{R}_1^S(F))$. Note that $a \in S$ but $b \notin S$, i.e., (F, a) is a yes-instance of $Cred_{naive,1}^{CPAF}$ but (F, b) is a no-instance of $Skept_{naive,1}^{CPAF}$.
- Assume φ is unsatisfiable. Consider some $S \subseteq A$. If there is some $x \in X$ such that neither $x \in S$ nor $\bar{x} \in S$, then $S \notin naive(\mathcal{R}_1^S(F))$. Likewise, if there is some $x \in X$ such that both $x \in S$ and $\bar{x} \in S$ then $S \notin cf(\mathcal{R}_1^S(F))$ and therefore $S \notin naive(\mathcal{R}_1^S(F))$. It remains to consider sets S in which for all $x \in X$ we have either $x \in S$ or $\bar{x} \in S$ but not both. Given such a set, we consider four cases:
 - (1) $a \notin S$ and $b \notin S$. Then $S \notin naive(\mathcal{R}_1^S(F))$ since $S \cup \{a\} \in cf(\mathcal{R}_1^S(F))$.
 - (2) $a \notin S$ and $b \in S$. Then $S \not\models \neg\varphi \wedge a \wedge \neg b$, i.e., $b \not\succeq_S a$. This means that the attack (a, b) is present in $\mathcal{R}_1^S(F)$. Note that every argument is either in S or in conflict with S in $\mathcal{R}_1^S(F)$. Moreover, $S \in cf(\mathcal{R}_1^S(F))$. We can conclude that $S \in naive(\mathcal{R}_1^S(F))$.




Fig. 7. Construction used in the proof of Lemma 19. Given a formula φ over variables $X = \{x_1, x_2, \dots, x_n\}$, a CPAF F is constructed such that φ is satisfiable iff (F, a) is a yes-instance of $Cred_{naive,1}^{CPAF}$ iff (F, b) is a no-instance of $Skept_{naive,1}^{CPAF}$.

- (3) $a \in S$ and $b \notin S$. Then $S \models \neg\varphi \wedge a \wedge \neg b$, i.e., $b \succ_S a$. This means that the attack (a, b) is deleted in $\mathcal{R}_1^S(F)$, which further implies that $S \cup \{b\} \in cf(\mathcal{R}_1^S(F))$. Thus, $S \notin naive(\mathcal{R}_1^S(F))$.
- (4) $a \in S$ and $b \in S$. Then $S \not\models \neg\varphi \wedge a \wedge \neg b$, i.e., $b \not\succeq_S a$. This means that the attack (a, b) is present in $\mathcal{R}_1^S(F)$. Thus, $S \notin cf(\mathcal{R}_1^S(F))$ and therefore $S \notin naive(\mathcal{R}_1^S(F))$.

In conclusion, if $S \in naive_{cp}^1(F)$ then $a \notin S$ and $b \in S$. Thus, (F, a) is a no-instance of $Cred_{naive,1}^{CPAF}$ but (F, b) is a yes-instance of $Skept_{naive,1}^{CPAF}$. \square

For naive semantics with global maximization, skeptical acceptance rises by even two levels in the polynomial hierarchy under Reduction 1. The reason for this is the increased complexity of the verification problem in this case (cf. Lemma 16). By QBF_{\forall}^2 we denote the Π_2^P -complete problem of deciding whether a quantified boolean formula of the form $\forall Y \exists Z \varphi$, where φ is a formula over $Y \cup Z$, is true.

Lemma 20. $Skept_{naive-glb,1}^{CPAF}$ is Π_2^P -complete.

Proof. Σ_2^P -membership for the complementary problem of $Skept_{naive-glb,1}^{CPAF}$: given a CPAF $F = (A, R, cond)$ and an argument x , guess a set $S \subset A$ and check that $x \notin S$ and, in coNP-time, that $S \in naive-glb_{cp}^1(F)$.

Π_2^P -hardness: let $\forall Y \exists Z \varphi$ be an arbitrary instance of QBF_{\forall}^2 over variables $Y = \{y_1, \dots, y_n\}$ and $Z = \{z_1, \dots, z_m\}$. Let $X = Y \cup Z$. Using fresh variables a and z_{m+1} we construct an instance (F, a) of $Skept_{naive-glb,1}^{CPAF}$ where $F = (A, R, cond)$ with

- $A = X \cup \{\bar{y} | y \in Y\} \cup \{a, z_{m+1}\}$,
- $R = \{(y, \bar{y}), (\bar{y}, y) | y \in Y\} \cup \{(z_i, z_j) | 1 \leq i < j \leq m + 1\} \cup \{(z_i, a) | 1 \leq i \leq m + 1\}$,
- and $cond$ defined by the following rules:
 - * $\varphi \wedge z_{m+1} \Rightarrow z_j \succ z_i$ for all $1 \leq i < j \leq m + 1$,
 - * $\varphi \wedge z_{m+1} \Rightarrow a \succ z_i$ for all $1 \leq i \leq m + 1$,
 - * $z_i \wedge \bigwedge_{1 \leq j \leq m+1, j \neq i} (\neg z_j) \Rightarrow a \succ z_i$ for all $1 \leq i \leq m$.

Expressed in natural language, the first two rules remove all conflicts between z_1, \dots, z_{m+1}, a if $\varphi \wedge z_{m+1}$ is satisfied, and the third rule removes the conflict between some $z \in Z$ and a if this z is the only element from Z that is part of the extension, and if z_{m+1} is also not part of the extension.

The above construction is visualized in Fig. 8. Note that the resulting CPAF is polynomial in the size of φ as we employ $O(m^2)$ rules, each linear in the size of φ . It remains to show that $\forall Y \exists Z \varphi$ is true iff (F, a) is a yes-instance of $Skept_{naive-glb,1}^{CPAF}$.

- Assume that $\forall Y \exists Z \varphi$ is true. We want to show that for all $S \in naive-glb_{cp}^1(F)$ we have $a \in S$. Towards a contradiction assume this is not the case, i.e., there is some $S \in naive-glb_{cp}^1(F)$ such that $a \notin S$. There are two possibilities:
 - (1) $S \models \varphi$. Then for $S' = S \cup \{a, z_{m+1}\}$ we also have $S' \models \varphi$ since a and z_{m+1} are fresh variables. Moreover, $S' \in cf_{cp}^1(F)$ since $S' \models \varphi \wedge z_{m+1}$ and thus all conflicts between the arguments z_1, \dots, z_{m+1}, a are removed. But $S \subset S'$, i.e., $S \notin naive-glb_{cp}^1(F)$. Contradiction.
 - (2) $S \not\models \varphi$. Then $S \not\models \varphi \wedge z_{m+1}$ and therefore the conflicts between z_1, \dots, z_{m+1}, a are not removed. This means that at most one of z_1, \dots, z_{m+1}, a is in S since we require $S \in cf_{cp}^1(F)$. Indeed, exactly one argument from z_1, \dots, z_{m+1}, a has to be in S , since if none of them were in S




Fig. 8. Construction used in the proof of Lemma 20. Given a quantified Boolean formula $\forall Y \exists Z \varphi$ over variables $Y = \{y_1, \dots, y_n\}$ and $Z = \{z_1, \dots, z_m\}$, a CPAF F is constructed such that $\forall Y \exists Z \varphi$ is true iff (F, a) is a yes-instance of $Skept_{naive-glb,1}^{CPAF}$.

then we could add any of these arguments to S and the resulting set would still be conflict free regardless of preferences. By our assumption, $a \notin S$. Again, we distinguish two cases:

- $z_i \in S$ with $1 \leq i \leq m$. But then $S \cup \{a\} \in cf_{cp}^1(F)$ because the following rule would apply: $z_i \wedge \bigwedge_{1 \leq j \leq m+1, j \neq i} (\neg z_j) \Rightarrow a \succ z_i$. Thus, $S \notin naive-glb_{cp}^1(F)$. Contradiction.
- $z_{m+1} \in S$. Let $I_Y = Y \cap S$. Since $\forall Y \exists Z \varphi$ is true there is some $I_Z \subseteq Z$ such that $I_Y \cup I_Z \models \varphi$. Therefore, $S' \models \varphi$ for $S' = I_Y \cup \{\bar{y} | y \notin I_Y\} \cup I_Z \cup \{a, z_{m+1}\}$. Moreover, $S' \in cf_{cp}^1(F)$ since $S' \models \varphi \wedge z_{m+1}$ and thus all conflicts between z_1, \dots, z_{m+1}, a are removed. Since $S \subset S'$ by construction we have that $S \notin naive-glb_{cp}^1(F)$. Contradiction.

In all cases we arrive at a contradiction, and we can conclude that $a \in S$ for all $S \in naive-glb_{cp}^1(F)$.

- Assume that $\forall Y \exists Z \varphi$ is not true. Then there is some $I_Y \subseteq Y$ such that $I_Y \cup I_Z \not\models \varphi$ for all $I_Z \subseteq Z$. Let $S = I_Y \cup \{\bar{y} | y \notin I_Y\} \cup \{z_{m+1}\}$. Clearly, $S \in cf_{cp}^1(F)$. Moreover, there can be no $S' \supset S$ such that $S' \in cf_{cp}^1(F)$ since we would need to add at least one argument from z_1, \dots, z_m, a to S . But these arguments are all in conflict with z_{m+1} unless $S' \models \varphi \wedge z_{m+1}$, which we know to be impossible. \square

We now turn our attention to admissibility-based semantics, where, in contrast to semantics based only on conflict-freeness, the choice of preference reduction makes no difference with regards to complexity. Again, let us first consider the cases in which there is no rise in complexity compared to AFs.

Lemma 21. $Cred_{\sigma,i}^{CPAF}$ is NP-complete for $\sigma \in \{adm, com, stb, prf-glb\}$ and $i \in \{1, 2, 3, 4\}$.

Proof. Hardness follows from hardness for AFs. Regarding membership of $\sigma \in \{adm, com, stb\}$, given a CPAF F and an argument x we can simply guess a set of arguments S containing x and, by Lemma 15, check whether $S \in \sigma_{cp}^i(F)$ in polynomial time. Regarding membership of $prf-glb$, it suffices to test whether (F, x) is a yes-instance of $Cred_{adm,i}^{CPAF}$. \square

Lemma 22. Let $i \in \{1, 2, 3, 4\}$. $Skept_{\sigma,i}^{CPAF}$ is trivial for $\sigma = adm$, coNP-complete for $\sigma = stb$, and Π_2^P -complete for $\sigma \in \{prf, prf-glb\}$.

Proof. Hardness follows from hardness for AFs. Regarding membership, let $F = (A, R, cond)$ be a CPAF and $x \in A$. Concerning $\sigma = adm$, note that \emptyset is always admissible in $\mathcal{R}_i^\emptyset(F)$, i.e., (F, x) is trivially a no-instance. Regarding $\sigma \in \{stb, prf, prf-glb\}$ we consider the complementary problem: guess a set $S \subset A$ and check that $x \notin S$ and that $S \in \sigma_{cp}^i(F)$ with $\sigma \in \{stb, prf, prf-glb\}$. Checking

$S \in \sigma_{cp}^i(F)$ can be done in polynomial time in the case of $\sigma = stb$ and in coNP-time in the case of $\sigma \in \{prf, prf-glb\}$ (cf. Lemma 15). \square

In the case of grounded semantics, both credulous and skeptical acceptance are located one level higher on the polynomial hierarchy compared to AFs. For complete semantics, the same is true for skeptical acceptance.

Lemma 23. *Let $i \in \{1, 2, 3, 4\}$. $Cred_{grd,i}^{CPAF}$ is NP-complete. $Skept_{\sigma,i}^{CPAF}$ is coNP-complete for $\sigma \in \{grd, com\}$.*

Proof. We will consider the complementary problems of $Skept_{grd,i}^{CPAF} / Skept_{com,i}^{CPAF}$ and show that they are NP-complete since this allows us to prove all results simultaneously.

NP-membership: given a CPAF $F = (A, R, cond)$ and an argument $x \in A$, guess a set $S \subseteq A$ and, in polynomial time, check whether $S \in grd_{cp}^i(F)$ and $x \in S$ (resp. $S \in grd_{cp}^i(F)$ and $x \notin S$ or $S \in com_{cp}^i(F)$ and $x \notin S$).

NP-hardness by reduction from SAT: let φ be an arbitrary propositional formula over variables X . Let a and b be fresh variables, i.e., $a, b \notin X$. We construct an instance (F, b) of $Cred_{grd,i}^{CPAF}$ as follows: $F = (A, R, cond)$ with $A = X \cup \{x^* | x \in X\} \cup \{a, b\}$, $R = \{(x, a), (a, x), (x, x^*), (x^*, x) | x \in X\} \cup \{(a, b)\}$, and $cond$ defined by the rules $\varphi \Rightarrow x \succ a$, $\neg\varphi \Rightarrow a \succ x$, and $x \Rightarrow x \succ x^*$ for $x \in X$, i.e., $x \succ_S a$ iff $S \models \varphi$, $a \succ_S x$ iff $S \models \neg\varphi$, and $x \succ_S x^*$ iff $S \models x$. Figure 9 depicts the above construction. In fact, this construction also works for the complementary problem of skeptical acceptance with respect to grounded and complete semantics, except that we will ask for the acceptance of the argument a instead of b . In this spirit, we now show that φ is satisfiable iff (F, b) is a yes-instance of $Cred_{grd,i}^{CPAF}$ iff (F, a) is a no-instance of $Skept_{grd,i}^{CPAF} / Skept_{com,i}^{CPAF}$.

- Suppose φ is satisfiable. Then there is an interpretation $I \subseteq X$ such that $I \models \varphi$. We assume that $I \neq \emptyset$. This is permissible since we can check in polynomial time whether \emptyset satisfies φ , and if this is the case, return a trivial yes-instance of $Cred_{grd,i}^{CPAF}$ (or a trivial no-instance of $Skept_{grd,i}^{CPAF} / Skept_{com,i}^{CPAF}$). Consider $S = I \cup \{b\}$. Then $S \models \varphi$ since $I \models \varphi$ and b does not occur in φ . We then have $x \succ_S a$ and $x \succ_S x^*$ for all $x \in I$, but $x \not\succ_S x^*$ for $x \in X \setminus I$. Thus, the unattacked arguments in $\mathcal{R}_i^S(F)$ are exactly those in I . Since $I \neq \emptyset$, b is defended in $\mathcal{R}_i^S(F)$ against a by the arguments in I . Thus, S is the minimal complete extension in $\mathcal{R}_i^S(F)$ and therefore also grounded in $\mathcal{R}_i^S(F)$. This implies that (F, b) is a yes-instance of $Cred_{grd,i}^{CPAF}$ while (F, a) is a no-instance of both $Skept_{grd,i}^{CPAF}$ and $Skept_{com,i}^{CPAF}$.
- Suppose φ is unsatisfiable. Then, for every $S \subseteq A$ and $x \in X$, we have $a \succ_S x$. Thus, the argument a is unattacked in every $\mathcal{R}_i^S(F)$, i.e., every complete extension (and therefore also every grounded extension) in F must contain a . Since a and b are in conflict, b is contained in no complete or




Fig. 9. Construction used in the proof of Lemma 23. Given a formula φ over variables $X = \{x_1, x_2, \dots, x_n\}$, a CPAF F is constructed such that φ is satisfiable iff (F, b) is a yes-instance of $Cred_{grd,i}^{CPAF}$ iff (F, a) is a no-instance of $Skept_{grd,i}^{CPAF} / Skept_{com,i}^{CPAF}$.

grounded extension. Thus, (F, b) is a no-instance of $Cred_{grd,i}^{CPAF}$ while (F, a) is a yes-instance of both $Skept_{grd,i}^{CPAF}$ and $Skept_{com,i}^{CPAF}$. \square

The last problem to consider is credulous acceptance for preferred semantics with local maximization. The following proof is the only one in this section to utilize one of the well-known standard reductions for AFs [29]. Only a very limited inclusion of conditional preferences is necessary. Indeed, only a single preference rule, consisting of a very simple propositional formula, is used in the construction.

Lemma 24. $Cred_{prf,i}^{CPAF}$ is Σ_2^P -complete for $i \in \{1, 2, 3, 4\}$.

Proof. Σ_2^P -membership: given a CPAF $F = (A, R, cond)$ and an argument x , guess a set $S \subset A$ and check that $x \in S$ and, in coNP-time, that $S \in prf_{cp}^i(F)$.

Π_2^P -hardness of the complementary problem: let $\forall X \exists Y \varphi$ be an arbitrary instance of QBF_{\forall}^2 in 3-CNF over variables $Y = \{y_1, \dots, y_n\}$ and $Z = \{z_1, \dots, z_m\}$ with clauses $C = \{c_1, \dots, c_k\}$. Let $X = Y \cup Z$. Using fresh variables a and b we construct an instance (F, a) of $co-Cred_{prf,i}^{CPAF}$ where $F = (A, R, cond)$,

- $A = X \cup \{\bar{x} \mid x \in X\} \cup C \cup \{a, b, \top, \perp\}$,
- $R = \{(x, \bar{x}), (\bar{x}, x) \mid x \in X\} \cup \{(x, c) \mid x \in C\} \cup \{(\bar{x}, c) \mid \neg x \in C\} \cup \{(c, c), (c, \top) \mid c \in C\} \cup \{(\perp, z), (\perp, \bar{z}) \mid z \in Z\} \cup \{(\top, \perp), (\perp, \perp)\} \cup \{(a, b), (b, a)\}$,
- and $cond$ defined by the following rule: $\bigvee_{z \in Z} (z \vee \bar{z}) \Rightarrow b \succ a$.

This construction is exemplified in Fig. 10. It remains to show that $\forall Y \exists Z \varphi$ is true iff $a \notin S$ for all $S \in prf_{cp}^i(F)$.

- Assume that $\forall Y \exists Z \varphi$ is true. Towards a contradiction, assume that there is some $S \in prf_{cp}^i(F)$ such that $a \in S$. Then it must be that $b \not\succeq_S a$, otherwise a is undefended in $\mathcal{R}_i^S(F)$. Thus, for all $z \in Z$ we have $z \notin S$ and $\bar{z} \notin S$. Let $I_Y = S \cap Y$. Since $\forall Y \exists Z \varphi$ is true there is some $I_Z \subseteq Z$ such that $I_Y \cup I_Z \models \varphi$. Let $S' = I_Y \cup \{\bar{y} \mid y \in Y, y \notin I_Y\} \cup I_Z \cup \{\bar{z} \mid z \in Z, z \notin I_Z\} \cup \{\top, a\}$. Clearly, $S \subset S'$. Moreover, S' is admissible in $\mathcal{R}_i^S(F)$: since $I_Y \cup I_Z \models \varphi$ all clause-arguments $c \in C$ are attacked by arguments in S' , and therefore \top is defended by S' . This further implies that all arguments z, \bar{z} are defended by S' against \perp . We can conclude that S is not preferred in $\mathcal{R}_i^S(F)$. Contradiction.




Fig. 10. Construction used in the proof of Lemma 24. Given the quantified Boolean formula $\forall y_1 y_2 \exists z_1 z_2 \varphi$, with φ consisting of clauses $c_1 = (y_1 \vee \neg y_2 \vee z_1)$, $c_2 = (\neg y_1 \vee \neg z_1 \vee z_2)$, and $c_3 = (y_2 \vee z_1 \vee \neg z_2)$, a CPAF F is constructed such that $\forall y_1 y_2 \exists z_1 z_2 \varphi$ is true iff (F, a) is a no-instance of $Cred_{prf,i}^{CPAF}$.

- Assume that $\forall Y \exists Z \varphi$ is not true. Then there is $I_Y \subseteq Y$ such that $I_Y \cup I_Z \not\models \varphi$ for all $I_Z \subseteq Z$. Let $S = I_Y \cup \{\bar{y} \mid y \in Y, y \notin I_Y\} \cup \{a\}$. Note that $b \not\prec_S a$ and that all arguments y, \bar{y} defend themselves. Thus, S is admissible in $\mathcal{R}_i^S(F)$. Towards a contradiction, assume there is $S' \supset S$ such that S' is admissible in $\mathcal{R}_i^S(F)$. This means one of the following must be the case:
 - * $\top \in S'$. Then \top needs to be defended by S' against the clause arguments $c \in C$. But this means that $I = (S' \cap Y) \cup (S' \cap Z)$ satisfies all clauses in φ , i.e., $I \models \varphi$. Note that S contains exactly one of y, \bar{y} for every $y \in Y$. Thus, $S' \cap Y = S \cap Y = I_Y$. The fact that $I_Y \cup (S' \cap Z) \models \varphi$ contradicts $I_Y \cup I_Z \not\models \varphi$ for all $I_Z \subseteq Z$.
 - * $z \in S'$ for some $z \in Z$. Then z needs to be defended by S' against \perp . This is only possible if $\top \in S'$, which we already have shown to not be the case. Contradiction.
 - * $\bar{z} \in S'$ for some $z \in Z$. Analogous to the case that $z \in S'$.

Since we arrive at a contradiction in all cases, $S \in \text{prf}_{cp}^i(F)$. Moreover, note that $a \in S$. \square

6. Related formalisms

We now investigate the connection between CPAFs and related formalisms. First, we show that Value-based Argumentation Frameworks (VAFs) [8,12] can be captured by CPAFs in a straightforward way. Secondly, we consider Extended Argumentation Frameworks (EAFs) [10,28,34] and highlight similarities and differences to CPAFs. Lastly, we compare our CPAFs with a recently introduced alternative approach to conditional preferences in abstract argumentation [2].

6.1. 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 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, $\text{val}: A \rightarrow V$ 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 $F_p = (A, R, V, \text{val}, \succ_p)$ 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 $\text{val}(y) \not\prec_p \text{val}(x)$. 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.

Figure 11a 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 Figs 11b and 11c.




Fig. 11. Example VAF with two audiences p_1 ($v_1 \succ v_2$) and p_2 ($v_2 \succ v_1$).




Fig. 12. CPAF obtained by translating the VAF of Fig. 11 according to Definition 17.

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 in $Tr(F)$ respectively.

Definition 17. Let $F = (A, R, V, val, P)$ be a VAF. Then $Tr(F) = (A', R', cond)$ is the CPAF such that

- $A' = A \cup P$,
- $R' = R \cup \{(p, p'), (p', p) | p, p' \in P, p \neq p'\}$,
- for every $S \subseteq A'$, $a \succ_S b$ iff there is $p \in P$ with $S \cap P = \{p\}$ and $val(a) \succ_p val(b)$.

Intuitively, each audience in the initial VAF is added as its own argument in our CPAF. The attacks of the VAF are preserved and symmetric attacks are added between all audience-arguments. Lastly, the preferences in our CPAF correspond to the preferences of each audience and are controlled by the newly introduced audience-arguments. Figure 12 shows the CPAF that results if the above translation is applied to the VAF of Fig. 11.

Observe that the successful attacks in some AVAF $F_p = (A, R, V, val, \succ_p)$ are also attacks in $\mathcal{R}_1^{S \cup \{p\}}(Tr(F))$, 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 25. 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 $S \cup \{p\} \in adm_{cp}^1(Tr(F))$.

Furthermore, 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_{cp}^1 -semantics. Therefore, the direct correspondence between admissible sets observed in Lemma 25 carries over to preferred extensions.

Proposition 26. 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.

It must be pointed out that the translation provided in Definition 17 was designed for VAFs in which each audience is given explicitly. However, VAFs can also be defined with preferences given implicitly as the set of all possible audiences [8], where each audience corresponds to a linear ordering over all values. In this case, the translation of Definition 17 is not polynomial as the number of audience arguments would be factorial in the number of values. We now provide an alternative translation that can handle this implicit definition of audiences and where we only need $|V|^2$ additional arguments.




Fig. 13. CPAF obtained by translating the VAF of Fig. 11 according to Definition 18.

Definition 18. Let $F = (A, R, V, val, P)$ be a VAF, with P implicitly given as the set of all possible linear orderings over V . Then $Tr_2(F) = (A', R', cond)$ is the CPAF such that

- $A' = A \cup \{v^k | v \in V, 1 \leq k \leq |V|\}$,
- $R' = R \cup \{(v^k, w^k), (w^k, v^k) | v \in V, w \in V, 1 \leq k \leq |V|, v \neq w\} \cup \{(v^k, v^l) | v \in V, k \neq l\}$,
- $cond$ is defined as follows: for every $a, b \in A$ such that $val(a) \neq val(b)$ and every $1 \leq k < |V|$ we introduce the rule $val(a)^k \wedge (\bigvee_{l=k+1}^{|V|} val(b)^l) \Rightarrow a > b$.

Figure 13 shows the CPAF that results if the above translation is applied to the VAF of Fig. 11. As with our first translation (cf. Definition 17), there is a direct semantic correspondence between the initial VAF and the constructed CPAF. The idea is the following: along with the arguments and attacks of the original VAF, we introduce arguments $v^1, \dots, v^{|V|}$ for each value $v \in V$. If v^k is accepted, this means that v is considered the k -th best value. Since v^k attacks all other value-arguments w^k with $w \neq v$ we know that no other value is simultaneously ascribed the k -th best position. Moreover, v^k attacks all v^l with $l \neq k$, i.e., v is only ascribed the k -th best position and no other. Then, we prefer an argument a to another argument b if the value of a is preferred (appears at an earlier position in the linear ordering) than b . In this way, each extension corresponds to a linear ordering over all values, i.e., each extension corresponds to an audience. This further implies that each S -reduct of the constructed CPAF has exactly the same attacks between the arguments of the initial VAFs as the AVAF corresponding to the value-ordering encoded in S . This gives us a result analogous to Proposition 26.

Proposition 27. Given a VAF $F = (A, R, V, val, P)$ where P is implicitly given as the set of all possible linear orderings over V , $x \in A$ is subjectively (resp. objectively) accepted in F iff x is credulously (resp. skeptically) preferred in $Tr_2(F)$ w.r.t. Reduction 1.

Our translations highlight 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.

Moreover, note that the problem of subjective acceptance in VAFs is NP-complete [12], even if the set of all audiences is represented implicitly. In contrast, we have shown that credulous acceptance in CPAFs is Σ_2^P -complete (cf. Table 4). Thus, assuming that the polynomial hierarchy does not collapse, finding a polynomial translation from CPAFs to VAFs analogous to our Proposition 26 (resp. Proposition 27) is not possible when considering credulously/subjectively accepted arguments.

6.2. 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 [34]. 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 Fig. 14b. Compare this to the formalization via a CPAF in Fig. 14a. 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 [34] states that admissibility for EAFs was defined in this way because it was deemed important to satisfy Dung's Fundamental Lemma [26], 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 of formalisms that can deal with conditional preferences in a flexible way. For example, in Fig. 14a 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. Fig. 2). In order to express our preference in case Person 1 has an alibi we extend our EAF from Fig. 14b by adding an attack from a_1 to the attack (c_1, c_2) , as shown in Fig. 15a.³ 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 Fig. 15b.




Fig. 14. A simplified version of the Sherlock Holmes example modeled via a CPAF and an EAF.

³The EAFs of Fig. 14b and Fig. 15a are also used as examples in Modgil's original paper [34].



Fig. 15. The full Sherlock Holmes example modeled by two different EAFs.

Thus, we can see that CPAFs allow for more flexibility when combining preferences associated with several arguments.

There are also some differences between CPAFs and EAFs when it comes to preferred semantics. For instance, stable extensions in EAFs are not necessarily preferred extensions [28]. In CPAFs, every stable extension is also preferred, except if we use global maximization and Reduction 1 (cf. Proposition 4). Moreover, credulous acceptance under preferred semantics is in NP for EAFs [28], but Σ_2^P -complete for CPAFs when using local maximization (cf. Table 4).

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 [35].

6.3. Lifting preferences over arguments to sets of arguments

In our CPAFs, we deal with preferences by using preference reductions which modify the attack relation (see Definition 5). There exist other approaches to preferences in argumentation, where preference orderings over arguments are lifted to sets of arguments [1,6,22,33], and the most preferred extensions are then selected according to this new preference ordering.

Recently [2], conditional preferences in abstract argumentation have been investigated using the aforementioned preference liftings. We refer to the CPAFs introduced in that work as *lifting-based* CPAFs. Similarly to our reduction-based CPAFs, a lifting-based CPAF is given as (A, R, Γ) where (A, R) is an AF and Γ is a set of conditional preference rules of the form $a_1 \succ a_2 \leftarrow b_1 \wedge \dots \wedge b_m \wedge \neg c_1 \wedge \dots \wedge \neg c_n$ built from arguments $a_1, a_2, b_1, \dots, b_m, c_1, \dots, c_n$. The conditional preferences over arguments given by Γ are lifted to preferences over sets of arguments according to one of three criteria (democratic, elitist, KTV), and then the ‘best’ extensions are selected according to this lifted preference ordering.

Note that lifting-based CPAFs, in contrast to our reduction-based CPAFs, satisfy principle $P2^*$ (cf. Definition 15) by design, since the ‘best’ extensions selected in a lifting-based CPAF (A, R, Γ) are always extensions of (A, R) . We note that, for complete and stable semantics, Reduction 3 satisfies $P2^*$ as well and thus selects extensions in the style of preference liftings (cf. Table 3).

The conditional preference rules Γ of a lifting-based CPAF (A, R, Γ) are usually assumed to be well-formed, which ensures that arguments a_1, a_2 occurring in the head of a rule $a_1 \succ a_2 \leftarrow b_1 \wedge \dots \wedge b_m \wedge \neg c_1 \wedge \dots \wedge \neg c_n$ do not occur in the body of the same rule. This is to prevent counterintuitive results, as explained in [2] via the following example: given (A, R, Γ) with extensions $\{\{a, b\}, \{a, c\}\}$ and Γ given

by $c \succ b \leftarrow b$ and $c \succ b \leftarrow c$, one would expect the only ‘best’ extension to be $\{a, c\}$. However, under semantics of lifting-based CPAFs, both $\{a, b\}$ and $\{a, c\}$ are ‘best’. This problem does not occur with the well-formed $\Gamma' = \{c \succ b \leftarrow\}$. In our reduction-based CPAFs we have no analogous assumption of well-formedness. Despite this, the counter-intuitive behavior observed above does not necessarily occur in our reduction-based CPAFs. For example, consider (A, R, cond) with $A = \{a, b, c\}$, $R = \{(b, c), (c, b)\}$, and cond given by the rules $b \Rightarrow c \succ b$ and $c \Rightarrow c \succ b$. Then $\text{prf}((A, R)) = \{\{a, b\}, \{a, c\}\}$. However, under all four preference reductions, the attack (b, c) is deleted as soon as b or c is in the extension under inspection. Thus, $\text{prf}_{cp}^i((A, R, \text{cond})) = \{\{a, c\}\}$.

Another difference between lifting-based CPAFs and our reduction-based CPAFs lies in their computational complexity, which is higher for lifting-based CPAFs in most cases. For example, verification for stable semantics is coNP-complete in lifting-based CPAFs [2] but remains in P in reduction-based CPAFs (see Table 4). As a result, credulous and skeptical acceptance for stable semantics are Σ_2^P -complete and Π_2^P -complete respectively in lifting-based CPAFs, while they remain NP-complete and coNP-completely respectively in reduction-based CPAFs. Some problems, such as credulous and skeptical acceptance of preferred semantics under elitist and KTV criteria, may even lie on the third level of the polynomial hierarchy for lifting-based CPAFs (tight bounds for the complexity of these problems have not been established yet). We observe that the increased complexity of lifting-based CPAFs is in many cases not due to the introduction of conditional preferences, but rather due to the preference-liftings themselves, as the complexity of lifting-based PAFs (featuring only unconditional preferences) is already considerably higher than that of regular AFs [1].

As pointed out in [2, Example 2], there are lifting-based CPAFs where not every (best) stable extension is also a (best) preferred extension. In contrast, every stable extension in a reduction-based CPAF is also a preferred extension, except when considering Reduction 1 and preferred semantics with global maximization (cf. Proposition 4).

To conclude this comparison, we want to emphasize that there is a conceptual difference between the reduction-based and lifting-based approaches to resolving preferences in argumentation: when using preference reductions, $x \succ y$ expresses that x is stronger than y ; when using preference liftings, $x \succ y$ expresses that we prefer outcomes containing x rather than y . Which of the two approaches should be chosen depends on the task at hand.

7. 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 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 satisfies many of the principles. We further investigate the computational complexity of CPAFs and show that this complexity can be influenced by the chosen preference reduction (in case of naive semantics) or by how maximization is handled (in case of naive and preferred semantics). Our results also 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).

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. We also discuss a

recently introduced alternative approach to conditional preferences in abstract argumentation, where preferences over arguments are lifted to preferences over sets of arguments.

The fact that Reduction 1 results in a higher complexity under naive semantics (cf. Table 4) is not unique to our setting. It has been shown that Reduction 1 causes a higher complexity compared to Reductions 2-4 also in the setting of claim-acceptance in AFs [13], although there the difference between the preference reductions extends to more than just naive semantics.

For future work, the relationship between CPAFs and existing approaches in structured argumentation [27] shall be investigated. Related to this point, it may also be interesting to see whether conditional preferences can be adapted to other formalisms such as bipolar argumentation frameworks [5], 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] or various forms of logic programming [20,21,23,24] relate to CPAFs.

Acknowledgements

We thank the anonymous reviewers, including the reviewers of the previous version of this paper, for their valuable feedback. This work was funded by the Austrian Science Fund (FWF) under grant P32830 and by the Vienna Science and Technology Fund (WWTF) under grant ICT19-065.

References

- [1] G. Alfano, S. Greco, F. Parisi and I. Trubitsyna, On preferences and priority rules in abstract argumentation, in: *Proc. IJCAI'22, IJCAI Organization*, 2022, pp. 2517–2524.
- [2] G. Alfano, S. Greco, F. Parisi, I. Trubitsyna et al., Abstract argumentation framework with conditional preferences, in: *Proc. AAI'23*, AAAI Press, 2023, pp. 6218–6227.
- [3] L. Amgoud and C. Cayrol, On the acceptability of arguments in preference-based argumentation, in: *Proc. UAI'98*, Morgan Kaufmann, 1998, pp. 1–7.
- [4] L. Amgoud and C. Cayrol, A reasoning model based on the production of acceptable arguments, *Ann. Math. Artif. Intell.* **34**(1–3) (2002), 197–215. doi:[10.1023/A:1014490210693](https://doi.org/10.1023/A:1014490210693).
- [5] L. Amgoud, C. Cayrol, M. Lagasque-Schiex and P. Livet, On bipolarity in argumentation frameworks, *Int. J. Intell. Syst.* **23**(10) (2008), 1062–1093. doi:[10.1002/int.20307](https://doi.org/10.1002/int.20307).
- [6] L. Amgoud and S. Vesic, Rich preference-based argumentation frameworks, *Int. J. Approx. Reason.* **55**(2) (2014), 585–606. doi:[10.1016/j.ijar.2013.10.010](https://doi.org/10.1016/j.ijar.2013.10.010).
- [7] S. Arora and B. Barak, *Computational Complexity: A Modern Approach*, Cambridge University Press, 2009.
- [8] K. Atkinson and T.J.M. Bench-Capon, Value-based argumentation, in: *Handbook of Formal Argumentation*, Vol. 2, College Publications, 2021, pp. 299–354. Also appears in *IfCoLog Journal of Logics and their Applications* 8(6), 1543–1588.
- [9] P. Baroni, M. Caminada and M. Giacomin, Abstract argumentation frameworks and their semantics, in: *Handbook of Formal Argumentation*, College Publications, 2018, pp. 159–236, Chapter 4.
- [10] P. Baroni, F. Cerutti, M. Giacomin and G. Guida, Encompassing attacks to attacks in abstract argumentation frameworks, in: *Proc. ECSQARU'09*, LNCS, Vol. 5590, Springer, 2009, pp. 83–94.
- [11] P. Baroni and M. Giacomin, On principle-based evaluation of extension-based argumentation semantics, *Artif. Intell.* **171**(10–15) (2007), 675–700. doi:[10.1016/j.artint.2007.04.004](https://doi.org/10.1016/j.artint.2007.04.004).
- [12] T.J.M. Bench-Capon, S. Doutre and P.E. Dunne, Audiences in argumentation frameworks, *Artif. Intell.* **171**(1) (2007), 42–71. doi:[10.1016/j.artint.2006.10.013](https://doi.org/10.1016/j.artint.2006.10.013).
- [13] M. Bernreiter, W. Dvořák, A. Rapberger and S. Woltran, The effect of preferences in abstract argumentation under a claim-centric view, in: *Proc. AAI'23*, AAAI Press, 2023, pp. 6253–6261.
- [14] M. Bernreiter, W. Dvořák and S. Woltran, Abstract argumentation with conditional preferences, in: *Proc. COMMA'22, Frontiers in Artificial Intelligence and Applications*, Vol. 353, IOS Press, 2022, pp. 92–103. doi:[10.3233/FAIA220144](https://doi.org/10.3233/FAIA220144).
- [15] M. Bernreiter, J. Maly and S. Woltran, Choice logics and their computational properties, *Artif. Intell.* **311** (2022), 103755. doi:[10.1016/j.artint.2022.103755](https://doi.org/10.1016/j.artint.2022.103755).

- [16] M. Bienvenu, J. Lang and N. Wilson, From preference logics to preference languages, and back, in: *Proc. KR'10*, AAAI Press, 2010, pp. 414–424.
- [17] S. Bistarelli and F. Santini, Weighted argumentation, in: *Handbook of Formal Argumentation*, Vol. 2, College Publications, 2021, pp. 355–395. Also appears in *IfCoLog Journal of Logics and their Applications* 8(6), 1589–1622.
- [18] J. Bobadilla, F. Ortega, A. Hernando and A. Gutiérrez, Recommender systems survey, *Knowl. Based Syst.* **46** (2013), 109–132. doi:[10.1016/j.knosys.2013.03.012](https://doi.org/10.1016/j.knosys.2013.03.012).
- [19] C. Boutilier, R.I. Brafman, C. Domshlak, H.H. Hoos and D. Poole, CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements, *J. Artif. Intell. Res.* **21** (2004), 135–191. doi:[10.1613/jair.1234](https://doi.org/10.1613/jair.1234).
- [20] G. Brewka, J.P. Delgrande, J. Romero and T. Schaub, Asprin: Customizing answer set preferences without a headache, in: *Proc. AAAI'15*, AAAI Press, 2015, pp. 1467–1474.
- [21] G. Brewka, I. Niemelä and T. Syrjänen, Logic programs with ordered disjunction, *Comput. Intell.* **20**(2) (2004), 335–357. doi:[10.1111/j.0824-7935.2004.00241.x](https://doi.org/10.1111/j.0824-7935.2004.00241.x).
- [22] G. Brewka, M. Truszczynski and S. Woltran, Representing preferences among sets, in: *Proc. AAAI'10*, AAAI Press, 2010, pp. 273–278.
- [23] A. Charalambidis, P. Rondogiannis and A. Troumpoukis, A logical characterization of the preferred models of logic programs with ordered disjunction, *Theory Pract. Log. Program.* **21**(5) (2021), 629–645. doi:[10.1017/S1471068421000235](https://doi.org/10.1017/S1471068421000235).
- [24] J.P. Delgrande, T. Schaub and H. Tompits, A framework for compiling preferences in logic programs, *Theory Pract. Log. Program.* **3**(2) (2003), 129–187. doi:[10.1017/S1471068402001539](https://doi.org/10.1017/S1471068402001539).
- [25] C. Domshlak, E. Hüllermeier, S. Kaci and H. Prade, Preferences in AI: An overview, *Artif. Intell.* **175**(7–8) (2011), 1037–1052. doi:[10.1016/j.artint.2011.03.004](https://doi.org/10.1016/j.artint.2011.03.004).
- [26] P.M. Dung, On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games, *Artif. Intell.* **77**(2) (1995), 321–358. doi:[10.1016/0004-3702\(94\)00041-X](https://doi.org/10.1016/0004-3702(94)00041-X).
- [27] P.M. Dung, P.M. Thang and T.C. Son, On structured argumentation with conditional preferences, in: *Proc. AAAI'19*, AAAI Press, 2019, pp. 2792–2800.
- [28] P.E. Dunne, S. Modgil and T.J.M. Bench-Capon, Computation in extended argumentation frameworks, in: *Proc. ECAI'10, FAIA*, Vol. 215, IOS Press, 2010, pp. 119–124.
- [29] W. Dvořák and P.E. Dunne, Computational problems in formal argumentation and their complexity, in: *Handbook of Formal Argumentation*, College Publications, 2018, pp. 631–687, Chapter 14. Also appears in *IfCoLog Journal of Logics and their Applications* 4(8), 2557–2622.
- [30] W. Dvořák, M. König, M. Ulbricht and S. Woltran, Rediscovering argumentation principles utilizing collective attacks, in: *Proc. KR'22*, 2022.
- [31] U. Endriss (ed.), *Trends in Computational Social Choice*, AI Access, 2017.
- [32] S. Kaci, L.W.N. van der Torre, S. Vesic and S. Villata, Preference in abstract argumentation, in: *Handbook of Formal Argumentation*, Vol. 2, College Publications, 2021, pp. 211–248.
- [33] S. Kaci, L.W.N. van der Torre and S. Villata, Preference in abstract argumentation, in: *Proc. COMMA'18, FAIA*, Vol. 305, IOS Press, 2018, pp. 405–412.
- [34] S. Modgil, Reasoning about preferences in argumentation frameworks, *Artif. Intell.* **173**(9–10) (2009), 901–934. doi:[10.1016/j.artint.2009.02.001](https://doi.org/10.1016/j.artint.2009.02.001).
- [35] S. Modgil and T.J.M. Bench-Capon, Metalevel argumentation, *J. Log. Comput.* **21**(6) (2011), 959–1003. doi:[10.1093/logcom/exq054](https://doi.org/10.1093/logcom/exq054).
- [36] S. Modgil and H. Prakken, Reasoning about preferences in structured extended argumentation frameworks, in: *Proc. COMMA'10, FAIA*, Vol. 216, IOS Press, 2010, pp. 347–358.
- [37] S. Modgil and H. Prakken, A general account of argumentation with preferences, *Artif. Intell.* **195** (2013), 361–397. doi:[10.1016/j.artint.2012.10.008](https://doi.org/10.1016/j.artint.2012.10.008).
- [38] S. Modgil and H. Prakken, Abstract rule-based argumentation, in: *Handbook of Formal Argumentation*, College Publications, 2018, pp. 287–364, Chapter 6. Also appears in *IfCoLog Journal of Logics and their Applications* 4(8), 2319–2406.
- [39] G. Pigozzi, A. Tsoukiàs and P. Viappiani, Preferences in artificial intelligence, *Ann. Math. Artif. Intell.* **77**(3–4) (2016), 361–401. doi:[10.1007/s10472-015-9475-5](https://doi.org/10.1007/s10472-015-9475-5).
- [40] Y. Shoham, Nonmonotonic logics: Meaning and utility, in: *Proc. IJCAI'87*, Morgan Kaufmann, 1987, pp. 388–393.
- [41] L. van der Torre and S. Vesic, The principle-based approach to abstract argumentation semantics, in: *Handbook of Formal Argumentation*, College Publications, 2018, pp. 797–838, Chapter 16. Also appears in *IfCoLog Journal of Logics and their Applications* 4(8), 2737–2780.