The Thirty-Seventh AAAI Conference on Artificial Intelligence (AAAI-23)

The Effect of Preferences in Abstract Argumentation under a Claim-Centric View

Michael Bernreiter, Wolfgang Dvořák, Anna Rapberger, Stefan Woltran

Institute of Logic and Computation, TU Wien, Austria
{mbernrei,dvorak,arapberg,woltran}@dbai.tuwien.ac.at

Abstract

In this paper, we study the effect of preferences in abstract argumentation under a claim-centric perspective. Recent work has revealed that semantical and computational properties can change when reasoning is performed on claim-level rather than on the argument-level, while under certain natural restrictions (arguments with the same claims have the same outgoing attacks) these properties are conserved. We now investigate these effects when, in addition, preferences have to be taken into account and consider four prominent reductions to handle preferences between arguments. As we shall see, these reductions give rise to different classes of claim-augmented argumentation frameworks, and behave differently in terms of semantic properties and computational complexity. This strengthens the view that the actual choice for handling preferences has to be taken with care.

1 Introduction

Arguments vary in their plausibility. Research in formal argumentation has taken up this aspect in both quantitative and qualitative terms (Li, Oren, and Norman 2011; Atkinson et al. 2017). Indeed, preferences are nowadays a standard feature of many structured argumentation formalisms (Modgil and Prakken 2013; Cyras and Toni 2016). At the same time, there are numerous generalizations of abstract Argumentation Frameworks (AFs) (Dung 1995) that consider the impact of preferences on the abstract level, be it in terms of argument strength (Kaci et al. 2021; Modgil 2009), preferences between values (Atkinson and Bench-Capon 2021), or weighted arguments/attacks (Bistarelli and Santini 2021). In AFs in which conflicts are expressed as a binary relation between arguments (*attack relation*), the incorporation of preferences typically results in the deletion or reversion of attacks. Deciding acceptability of arguments via argumentation semantics is thus reflected in terms of the modified attack relation (Kaci et al. 2021).

The difference in argument strength and the resulting modification of the attack relation naturally influences the acceptability of the arguments' conclusion (the *claim* of the argument). Claim acceptance in argumentation systems, i.e., the evaluation of commonly acceptable statements while disregarding their particular justifications, is an integral part of many structured argumentation formalisms (Modgil and Prakken 2018; Dung, Kowalski, and Toni 2009) and has received increasing attention in the literature (Horty 2002; Baroni and Riveret 2019; Dvořák and Woltran 2020; Rocha and Cozman 2022). A simple yet powerful generalization of AFs that allow for claim-based evaluation are Claim-augmented AFs (CAFs) (Dvořák and Woltran 2020), where each argument is assigned a claim. Semantics for CAFs can be obtained by evaluating the underlying AF before inspecting the claims of the acceptable arguments in the final step. CAFs serve as an ideal target formalism for ASPIC+ (Modgil and Prakken 2018) and other formalisms which utilize abstract argumentation semantics whilst also considering the claims of the arguments in the evaluation. Moreover, CAF semantics capture semantics of logic programs without the need of additional mappings (Rapberger 2020), in contrast to classical AF-instantiations (Caminada et al. 2015). Thus, we obtain a direct correspondence between claim-extensions in the CAF and conclusion-extensions in the original formalism.

Although the acceptance of claims is closely related to argument acceptance, there are subtle differences as observed in (Dvořák and Woltran 2020; Prakken and Vreeswijk 2002; Modgil and Prakken 2018) stemming from the fact that claims can appear as conclusion of several different arguments. As a consequence, several properties of AF semantics cannot be taken for granted when considered in terms of the arguments' claims. For instance, the property of *I-maximality*, i.e., $\subseteq$-maximality of extensions, which gives insights into the expressiveness of semantics (Dunne et al. 2015) and skeptical argument justification (Baroni and Giacomin 2007) is not satisfied by most CAF semantics (Dvořák, Rapberger, and Woltran 2020). Furthermore, the additional level of claims causes a rise in the computational complexity of standard decision problems (in particular, verification is one level higher in the polynomial hierarchy as for standard AFs), see (Dvořák et al. 2021). Luckily, these drawbacks can be alleviated by taking fundamental properties of the attack relation into account: the basic observation that attacks typically depend on the claim of the attacking arguments gives rise to the central class of *well-formed CAFs*. This class satisfies that all arguments with the same claim attack the same arguments; thus modeling a very natural behavior of arguments that is common to all leading structured argumentation formalisms and instantiations. Well-formed CAFs have the main advantage that most of the semantics behave 'as expected', e.g., they retain I-maximality, and their computational complexity is located at the same level of the polynomial hierarchy as for AFs.

Unfortunately, it turns out that well-formedness cannot be assumed if one deals with preferences in argumentation, as arguments with the same claim are not necessarily equally plausible. The following example demonstrates this.

Example 1. *Consider two arguments a, a' with claim α, and another argument b having claim β. Moreover, both a and a' attack b, while b attacks a. Furthermore assume that we are given the additional information that b is preferred over a' (for example, if assumptions in the support of b are stronger than assumptions made by a'). A common method to integrate such information on argument rankings is to* delete *attacks from arguments that attack preferred arguments. In this case, we delete the attack from a' to b.*

Both frameworks are depicted below: F represents the original situation while F' is the CAF resulting from deleting the unsuccessful attack from a' on the argument b.

Note that F is well-formed since all arguments with the same claims attack the same arguments. The unique acceptable argument-set w.r.t. stable semantics (cf. Definition 2) is $\{a, a'\}$ which translates to $\{\alpha\}$ on the claim-level.

The CAF F', on the other hand, is no longer well-formed since a' does not attack b. In F', the argument-sets $\{a, a'\}$ and $\{a', b\}$ are both acceptable w.r.t. to stable semantics. In terms of claims this translates to $\{\alpha\}$ and $\{\alpha, \beta\}$, which shows that I-maximality is violated on the claim-level.

Although well-formedness can not be guaranteed in view of preferences, this does not imply arbitrary behavior of the resulting CAF: on the one hand, preferences conform to a certain type of ordering (e.g., asymmetric, transitive) over the set of arguments; on the other hand, it is evident that the deletion, reversion, and other types of attack manipulation impose restrictions on the structure of the resulting CAF. Combining both aspects, we obtain that, assuming well-formedness of the initial framework, it is unlikely that preference incorporation results in arbitrary behavior. The key motivation of this paper is to identify and exploit structural properties of preferential argumentation in the scope of claim acceptance. The aforementioned restrictions suggest beneficial impact on both the computational complexity and on desired semantical properties such as I-maximality.

In this paper, we tackle this issue by considering four commonly used methods, so-called reductions, to integrate preference orderings into the attack relation: the most common modification is the deletion of attacks in case the attacking argument is less preferred than its target. This method is typically utilized to transform preference-based argumentation frameworks (PAFs) (Amgoud and Cayrol 1998) into AFs but is also used in many structured argumentation formalisms such as ASPIC+. This reduction has been criticized due to several problematic side-effects, e.g., it can be the case that two conflicting arguments are jointly acceptable, and has been accordingly adapted in (Amgoud and Vesic 2014); two other reductions have been introduced in (Kaci, van der Torre, and Villata 2018). We apply these four preference reductions to well-formed CAFs. In particular, our main contributions are as follows:

- For each of the four reductions, we characterize the possible structure of CAFs that are obtained by applying the reduction to a well-formed CAF and a preference relation. This results in four novel CAF classes, each of which constitutes a proper extension of well-formed CAFs not retaining full expressiveness of general CAFs. We investigate the relationship between these classes.
- We study I-maximality of stable, preferred, semi-stable, stage, and naive semantics of the novel CAF classes. Our results highlight a significant advantage of a particular reduction: we show that, for admissible-based semantics, this modification preserves I-maximality. The other reductions fail to preserve I-maximality; moreover, for naive and stage semantics, I-maximality cannot be guaranteed for any of the four reductions.
- Finally, we investigate the complexity of reasoning for CAFs with preferences with respect to conflict-free, admissible, complete, and all of the aforementioned semantics. We show that for three of the four reductions, the verification problem drops by one level in the polynomial hierarchy for all except complete semantics and is thus not harder than for well-formed CAFs (which in turn has the same complexity as the corresponding AF problems). Complete semantics remain hard for all but one preference reduction. Moreover, it turns out that verification for the reduction which deletes attacks from weaker arguments remains as hard as for general CAFs.

Our results constitute a systematic study of the structural and computational effect of preferences on claim acceptance. Since we use CAFs as our base formalism, our investigations extend to large classes of formalisms that can be represented as CAFs, just like results on AFs yield insights for formalisms that can be captured by AFs.

This paper is organized as follows. In Section 2, we recall necessary background. In Section 3, we introduce preference-based CAFs (PCAFs) which combine PAFs with well-formed CAFs. We characterize the novel CAF classes based on the preference reductions in Section 4, study the I-maximality of the semantics in Section 5, and their computational complexity in Section 6. We conclude in Section 7.

2 Preliminaries

We first define (abstract) argumentation frameworks (Dung 1995). U denotes a countable infinite domain of arguments.

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

Given an AF $F = (A, R)$ it can be convenient to write $a \in F$ for $a \in A$ and $(a, b) \in F$ for $(a, b) \in R$. 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 (Baroni, Caminada, and Giacomin 2018). We consider for σ the functions cf (conflict-free), adm (admissible), com (complete), $naive$ (naive), stb (stable), prf (preferred), sem (semi-stable), and stg (stage).

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

- *$S \in adm(F)$ 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;*
- *$S \in naive(F)$ iff there is no $T \in cf(F)$ with $S \subset T$;*
- *$S \in stb(F)$ iff each $a \in A \setminus S$ is attacked by S in F;*
- *$S \in prf(F)$ iff $S \in adm(F)$ and there is no $T \in adm(F)$ with $S \subset T$;*
- *$S \in sem(F)$ iff $S \in adm(F)$ and there is no $T \in adm(F)$ with $S_F^{\oplus} \subset T_F^{\oplus}$;*
- *$S \in stg(F)$ iff there is no $T \in cf(F)$ with $S_F^{\oplus} \subset T_F^{\oplus}$.*

Example 2. *Consider the AF $F = (\{a, a', b\}, \{(a, b), (a', b), (b, a)\})$ from Example 1, ignoring claims α and β. Then $cf(F) = \{\emptyset, \{a\}, \{a'\}, \{b\}, \{a, a'\}\}$, $adm(F) = \{\emptyset, \{a\}, \{a'\}, \{a, a'\}\}$, $naive(F) = \{\{b\}, \{a, a'\}\}$, and $\sigma(F) = \{\{a, a'\}\}$ for $\sigma \in \{com, stb, prf, sem, stg\}$.*

CAFs generalize AFs by assigning each argument a claim (Dvořák and Woltran 2020). We fix a countable infinite domain of claims $\mathcal{C}$.

Definition 3. *A claim-augmented argumentation framework (CAF) is a triple (A, R, cl) where (A, R) is an AF and $cl\colon A \to \mathcal{C}$ is a function that maps arguments to claims. The claim-function is extended to sets of arguments via $cl(E) = \{cl(a) \mid a \in E\}$. A well-formed CAF (wfCAF) is a CAF (A, R, cl) in which all arguments with the same claim attack the same arguments, i.e., for all $a, b \in A$ with $cl(a) = cl(b)$ we have $\{c \mid (a, c) \in R\} = \{c \mid (b, c) \in R\}$.*

The semantics of CAFs are based on those of AFs.

Definition 4. *Let $F = (A, R, cl)$ be a CAF. The claim-based variant of a semantics σ is defined as $\sigma_c(F) = \{cl(S) \mid S \in \sigma((A, R))\}$.*

Example 3. *Consider the CAF F from Example 1. Formally, $F = (A, R, cl)$ with $A = \{a, a', b\}$, $R = \{(a, b), (a', b), (b, a)\}$, $cl(a) = cl(a') = \alpha$, and $cl(b) = \beta$. F is well-formed and the underlying AF of F was investigated in Example 2. From there we can infer that, e.g., $cf_c(F) = \{\emptyset, \{\alpha\}, \{\beta\}\}$, $adm_c(F) = \{\emptyset, \{\alpha\}\}$, $naive_c(F) = \{\{\alpha\}, \{\beta\}\}$, and $stb_c(F) = \{\{\alpha\}\}$.*

Well-known basic relations between different AF semantics σ also hold for σ_c: $stb_c(F) \subseteq sem_c(F) \subseteq prf_c(F) \subseteq adm_c(F)$ as well as $stb_c(F) \subseteq stg_c(F) \subseteq naive_c(F) \subseteq cf_c(F)$ (Dvořák, Rapberger, and Woltran 2020).

Note that the semantics $\sigma \in \{naive, stb, prf, sem, stg\}$ employ argument maximization and result in incomparable extensions on regular AFs: for all $S, T \in \sigma(F)$, $S \subseteq T$ implies $S = T$. This property is called I-maximality (Baroni and Giacomin 2007), and is defined analogously for CAFs:

	$naive_c$	stb_c	prf_c	sem_c	stg_c
CAF	x	x	x	x	x
$wfCAF$	x	✓	✓	✓	✓

Table 1: I-maximality of CAFs.

σ	$Cred_\sigma^\Delta$	$Skept_\sigma^\Delta$	Ver_σ^{CAF}	Ver_σ^{wfCAF}
cf	in P	trivial	NP-c	in P
adm	NP-c	trivial	NP-c	in P
com	NP-c	P-c	NP-c	in P
$naive$	in P	coNP-c	NP-c	in P
stb	NP-c	coNP-c	NP-c	in P
prf	NP-c	Π_2^P-c	Σ_2^P-c	coNP-c
sem/stg	Σ_2^P-c	Π_2^P-c	Σ_2^P-c	coNP-c

Table 2: Complexity of CAFs ($\Delta \in \{CAF, wfCAF\}$).

Definition 5. *σ_c is I-maximal for a class $\mathcal{F}$ of CAFs if, for all CAFs $F \in \mathcal{F}$ and all $S, T \in \sigma_c(F)$, $S \subseteq T$ implies $S = T$.*

Table 1 shows I-maximality properties of CAFs (Dvořák, Rapberger, and Woltran 2020), revealing an important property of wfCAFs compared to general CAFs: I-maximality is preserved in all semantics except $naive_c$, implying natural behavior of these maximization-based semantics analogous to regular AFs; see, e.g., (van der Torre and Vesic 2017) for a general discussion of such properties.

Regarding computational complexity, we consider the following decision problems pertaining to CAF-semantics σ_c:

- *Credulous Acceptance* ($Cred_\sigma^{CAF}$): Given a CAF F and claim α, is α contained in some $S \in \sigma_c(F)$?
- *Skeptical Acceptance* ($Skept_\sigma^{CAF}$): Given a CAF F and claim α, is α contained in each $S \in \sigma_c(F)$?
- *Verification* (Ver_σ^{CAF}): Given a CAF F and a set of claims S, is $S \in \sigma_c(F)$?

We furthermore consider these reasoning problems restricted to wfCAFs and denote them by $Cred_\sigma^{wfCAF}$, $Skept_\sigma^{wfCAF}$, and Ver_σ^{wfCAF}. Table 2 shows the complexity of these problems (Dvořák and Woltran 2020; Dvořák et al. 2021). Here we see that the complexity of the verification problem drops by one level in the polynomial hierarchy when comparing general CAFs to wfCAFs. This is an important advantage of wfCAFs, as a lower complexity in the verification problem allows for a more efficient enumeration of claim-extensions (cf. (Dvořák and Woltran 2020)).

3 Preference-based CAFs

As discussed in the previous sections, wfCAFs are a natural subclass of CAFs with advantageous properties in terms of I-maximality and computational complexity. However, when resolving preferences among arguments, the resulting CAFs

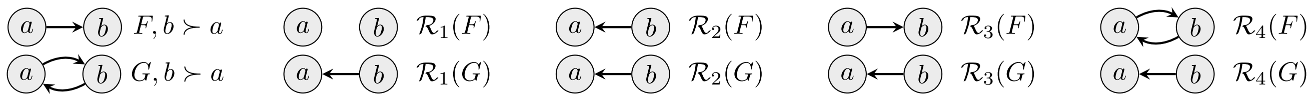

Figure 1: Effect of the four reductions on the attack relation between two arguments.

are typically no longer well-formed (cf. Example 1). In order to study preferences under a claim-centric view we introduce preference-based CAFs. These frameworks enrich the notion of wfCAFs with the concept of argument strength in terms of preferences. Our main goals are then to understand the effect of resolved preferences on the structure of the underlying wfCAF on the one hand, and to determine whether the advantages of wfCAFs are maintained on the other hand. Given this motivation, it is reasonable to consider the impact of preferences on *well-formed* CAFs only.

Definition 6. *A preference-based claim-augmented argumentation framework (PCAF) is a quadruple $F = (A, R, cl, \succ)$ where (A, R, cl) is a well-formed CAF and $\succ$ is an asymmetric preference relation over A.*

Notice that preferences in PCAFs are not required to be transitive. While transitivity of preferences is often assumed in argumentation (Amgoud and Vesic 2014; Kaci, van der Torre, and Villata 2018), it cannot always be guaranteed in practice (Kaci et al. 2021). However, we will consider the effect of transitive orderings when applicable.

If a and b are arguments and $a \succ b$ holds then we say that a is stronger than b. But what effect should this ordering have? How should this influence, e.g., the set of admissible arguments? One possibility is to remove all attacks from weaker to stronger arguments in our PCAF, and to then determine the set of admissible arguments in the resulting CAF. This altering of attacks in a PCAF based on its preference-ordering is called a reduction. The literature describes four such reductions for regular AFs (Amgoud and Cayrol 2002; Amgoud and Vesic 2014; Kaci, van der Torre, and Villata 2018), which we now adapt.

Definition 7. *Given a PCAF $F = (A, R, cl, \succ)$, the corresponding CAF $\mathcal{R}_i(F) = (A, R', cl)$ is constructed via Reduction i, where $i \in \{1, 2, 3, 4\}$, as follows:*

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

Figure 1 visualizes the above reductions. Intuitively, Reduction 1 removes attacks that contradict the preference ordering while Reduction 2 reverts such attacks. Reduction 3 removes attacks that contradict the preference ordering, but only if the weaker argument is attacked by the stronger argument also. Reduction 4 can be seen as a combination of Reductions 2 and 3. Observe that all four reductions are polynomial time computable with respect to the input PCAF. Note that many structured argumentation formalisms use preference reductions. For instance, ABA+ (Cyras and Toni 2016) employs attack reversal similar to Reduction 2 while some instances of ASPIC (Modgil and Prakken 2013) delete attacks from weaker arguments in the spirit of Reduction 1.

The semantics for PCAFs can now be defined in a straightforward way: first, one of the four reductions is applied to the given PCAF; then, CAF-semantics are applied to the resulting CAF.

Definition 8. *Let F be a PCAF and let $i \in \{1, 2, 3, 4\}$. The preference-claim-based variant of a semantics σ relative to Reduction i is defined as $\sigma_p^i(F) = \sigma_c(\mathcal{R}_i(F))$.*

Example 4. *Let $F = (A, R, cl, \succ)$ be the PCAF where $A = \{a, a', b\}$, $R = \{(a, b), (a', b), (b, a)\}$, $cl(a) = cl(a') = \alpha$, $cl(b) = \beta$, and $b \succ a'$. The underlying CAF (A, R, cl) of F was examined in Example 3.*

$\mathcal{R}_1(F) = (A, R', cl)$ with $R' = \{(a, b), (b, a)\}$, which is the same CAF as F' in Example 1. It can be verified that, e.g., $adm_p^1(F) = adm_c(\mathcal{R}_1(F)) = \{\{\emptyset, \{\alpha\}, \{\beta\}, \{\alpha, \beta\}\}$ and $stb_p^1(F) = \{\{\alpha\}, \{\alpha, \beta\}\}$.

Indeed, the choice of reduction can influence the extensions of a PCAF. For example, $\mathcal{R}_2(F) = (A, R'', cl)$ with $R'' = \{(a, b), (b, a), (b, a)\}$, $adm_p^2(F) = \{\emptyset, \{\alpha\}, \{\beta\}\}$, and $stb_p^2(F) = \{\{\alpha\}, \{\beta\}\}$.

It is easy to see that basic relations between semantics carry over from CAFs, as, if we have $\sigma_c(F) \subseteq \tau_c(F)$ for two semantics σ, τ and all CAFs F, then also $\sigma_p^i(F) \subseteq \tau_p^i(F)$ for all PCAFs F. It thus holds that for all $i \in \{1, 2, 3, 4\}$, $stb_p^i(F) \subseteq sem_p^i(F) \subseteq prf_p^i(F) \subseteq adm_p^i(F)$ as well as $stb_p^i(F) \subseteq stg_p^i(F) \subseteq naive_p^i(F) \subseteq cf_p^i(F)$.

Remark. *In this paper we require the underlying CAF of a PCAF to be well-formed. The reason for this is that we are interested in whether the benefits of well-formed CAFs are preserved when preferences have to be taken into account. Even from a technical perspective, admitting PCAFs with a non-well-formed underlying CAF is not very interesting with respect to the questions addressed in this paper. Indeed, any CAF could be obtained from such general PCAFs, regardless of which preference reduction we are using, by simply specifying the desired CAF and an empty preference relation. Thus, such general PCAFs have the same properties regarding I-maximality and complexity as general CAFs.*

4 Characterization & Expressiveness

Our first step towards understanding the effect of preferences on wfCAFs is to examine the impact of resolving preferences on the structure of the underlying CAF. To this end, we consider four new CAF classes which are obtained from applying the reductions of Definition 7 to PCAFs.

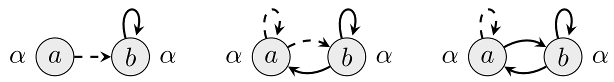

Figure 2: CAFs contained only in $\mathcal{R}_1$-$\mathbf{CAF}$, $\mathcal{R}_2$-$\mathbf{CAF}$, and $\mathcal{R}_4$-$\mathbf{CAF}$ respectively. Solid arrows are attacks, dashed arrows indicate missing attacks for the CAF to be well-formed.

Definition 9. *$\mathcal{R}_i$-$\mathbf{CAF}$ denotes the set of CAFs that can be obtained by applying Reduction i to PCAFs, i.e., $\mathcal{R}_i$-$\mathbf{CAF} = \{\mathcal{R}_i(F) \mid F$ is a PCAF$\}$.*

It is easy to see that $\mathcal{R}_i$-$\mathbf{CAF}$, with $i \in \{1, 2, 3, 4\}$, contains all wfCAFs (we can simply specify the desired wfCAF and an empty preference relation). However, not all CAFs are contained in $\mathcal{R}_i$-$\mathbf{CAF}$. For example, $F = (\{a, b\}, \{(a, b), (b, a)\}, cl)$ with $cl(a) = cl(b)$ can not be obtained from a PCAF F': such F' would need to contain either (a, b) or (b, a). But then, since the underlying CAF of a PCAF must be well-formed, F' would have to contain a self-attack which can not be removed by any of the reductions. This is enough to conclude that the four new classes are located in-between wfCAFs and general CAFs:

Proposition 1. *Let $\mathbf{CAF}$ be the set of all CAFs and $\mathbf{wfCAF}$ the set of all wfCAFs. For all $i \in \{1, 2, 3, 4\}$ it holds that $\mathbf{wfCAF} \subset \mathcal{R}_i$-$\mathbf{CAF} \subset \mathbf{CAF}$.*

Furthermore, the new classes are all distinct from each other, i.e., we are indeed dealing with *four* new CAF classes. Examples for CAFs that are in only one of $\mathcal{R}_1$-$\mathbf{CAF}$, $\mathcal{R}_2$-$\mathbf{CAF}$, and $\mathcal{R}_4$-$\mathbf{CAF}$ are provided in Figure 2. Moreover, $\mathcal{R}_3$-$\mathbf{CAF}$ is contained in all other three classes since Reduction 3 only alters symmetric attacks (in which case all reductions behave in the same way).

Proposition 2. *For all $i \in \{1, 2, 4\}$ and all $j \in \{1, 2, 3, 4\}$ such that $i \neq j$ it holds that $\mathcal{R}_i$-$\mathbf{CAF} \not\subseteq \mathcal{R}_j$-$\mathbf{CAF}$ and $\mathcal{R}_3$-$\mathbf{CAF} \subset \mathcal{R}_i$-$\mathbf{CAF}$.*

While the classes $\mathcal{R}_1$-$\mathbf{CAF}$, $\mathcal{R}_2$-$\mathbf{CAF}$, and $\mathcal{R}_4$-$\mathbf{CAF}$ are incomparable, we observe $\mathcal{R}_3$-$\mathbf{CAF} \subset \mathcal{R}_i$-$\mathbf{CAF}$ which reflects that Reduction 3 is the most conservative of the four reductions, removing attacks only when there is a counter-attack from the stronger argument.

We now know that applying preferences to wfCAFs results in four distinct CAF-classes that lie in-between wfCAFs and general CAFs. It is still unclear, however, how to determine whether some CAF belongs to one of these classes or not. Especially for $\mathcal{R}_2$-$\mathbf{CAF}$ and $\mathcal{R}_4$-$\mathbf{CAF}$ this is not straightforward, since Reductions 2 and 4 not only remove but also introduce attacks and therefore allow for many possibilities to obtain a particular CAF as result. We tackle this problem by characterizing the new classes via the so-called wf-problematic part of a CAF.

Definition 10. *A pair of arguments (a, b) is wf-problematic in a CAF $F = (A, R, cl)$ if $a, b \in A$, $(a, b) \notin R$, and there is $a' \in A$ with $cl(a') = cl(a)$ and $(a', b) \in R$. The set $wfp(F) = \{(a, b) \mid (a, b)$ is wf-problematic in $F\}$ is called the wf-problematic part of F.*

Intuitively, the wf-problematic part of a CAF F consists of those attacks that are missing for F to be well-formed (cf. Figure 2). Indeed, F is a wfCAF if and only if $wfp(F) = \emptyset$. The four new classes can be characterized as follows:

Proposition 3. *Let $F = (A, R, cl)$ be a CAF. Then*

- *$F \in \mathcal{R}_1$-$\mathbf{CAF}$ iff $(a, b) \in wfp(F)$ implies $(b, a) \notin wfp(F)$;*
- *$F \in \mathcal{R}_2$-$\mathbf{CAF}$ iff there are no arguments a, a', b, b' in F with $cl(a) = cl(a')$ and $cl(b) = cl(b')$ such that $(a, b) \in wfp(F)$, $(b, a) \notin R$, $(a', b) \in R$, and either $(b, a') \in R$ or $((a', b') \notin R$ and $(b', a') \notin R)$;*
- *$F \in \mathcal{R}_3$-$\mathbf{CAF}$ iff $(a, b) \in wfp(F)$ implies $(b, a) \in R$;*
- *$F \in \mathcal{R}_4$-$\mathbf{CAF}$ iff there are no arguments a, a', b, b' in F with $cl(a) = cl(a')$ and $cl(b) = cl(b')$ such that $(a, b) \in wfp(F)$, $(b, a) \notin R$, $(a', b) \in R$, and either $(b, a') \notin R$ or $((a', b') \notin R$ and $(b', a') \notin R)$.*

The above characterizations give us some insights into the effect of the various reductions on wfCAFs. Indeed, the similarity between the characterizations of $\mathcal{R}_1$-$\mathbf{CAF}$ and $\mathcal{R}_3$-$\mathbf{CAF}$, resp. $\mathcal{R}_2$-$\mathbf{CAF}$ and $\mathcal{R}_4$-$\mathbf{CAF}$, can intuitively be explained by the fact that Reductions 1 and 3 only remove attacks, while Reductions 2 and 4 can also introduce attacks. Furthermore, Proposition 3 allows us to decide in polynomial time whether a given CAF F can be obtained by applying one of the four preference reductions to a PCAF.

But what happens if we restrict ourselves to transitive preferences? Analogously to $\mathcal{R}_i$-$\mathbf{CAF}$, by $\mathcal{R}_i$-$\mathbf{CAF}_{tr}$ we denote the set of CAFs obtained by applying Reduction i to PCAFs with a transitive preference relation. It is clear that $\mathcal{R}_i$-$\mathbf{CAF}_{tr} \subseteq \mathcal{R}_i$-$\mathbf{CAF}$ for all $i \in \{1, 2, 3, 4\}$. Interestingly, the relationship between the classes $\mathcal{R}_i$-$\mathbf{CAF}_{tr}$ is different to that between $\mathcal{R}_i$-$\mathbf{CAF}$ (Proposition 2). Specifically, $\mathcal{R}_3$-$\mathbf{CAF}_{tr}$ is not contained in the other classes. Intuitively, this is because, in certain PCAFs F, transitivity can force $a_1 \succ a_n$ via $a_1 \succ a_2 \succ \ldots \succ a_n$ such that $(a_n, a_1) \in F$ but $(a_1, a_n) \notin F$. In this case, only Reduction 3 leaves the attacks between a_1 and a_n unchanged.

Proposition 4. *For all $i, j \in \{1, 2, 3, 4\}$ such that $i \neq j$ it holds that $\mathcal{R}_i$-$\mathbf{CAF}_{tr} \not\subseteq \mathcal{R}_j$-$\mathbf{CAF}_{tr}$.*

We will not characterize all four classes $\mathcal{R}_i$-$\mathbf{CAF}_{tr}$. However, capturing $\mathcal{R}_1$-$\mathbf{CAF}_{tr}$ will prove useful when analyzing the computational complexity of PCAFs using Reduction 1. Note that $wfp(F)$ can be seen as a directed graph, with an edge between vertices a and b whenever $(a, b) \in wfp(F)$. Thus, we may use notions such as paths and cycles in the wf-problematic part of a CAF.

Proposition 5. *$F \in \mathcal{R}_1$-$\mathbf{CAF}_{tr}$ for a CAF F iff (1) $wfp(F)$ is acyclic and (2) $(a, b) \in F$ implies that there is no path from a to b in $wfp(F)$.*

From the high-level point of view, our characterization results yield insights into the expressiveness of argumentation formalisms that allow for preferences. Propositions 3 and 5 show which situations can be captured by formalisms which (i) constructs attacks based on the claim of the attacking argument (i.e., formalisms with well-formed attack relation) and (ii) incorporate asymmetric or transitive preference relations on arguments using one of the four reductions.

	$naive_p^i$	stb_p^i	prf_p^i	sem_p^i	stg_p^i
$i \in \{1,2,4\}$	x	x	x	x	x
$i=3$	x	✓	✓	✓	x

Table 3: I-maximality of PCAFs.

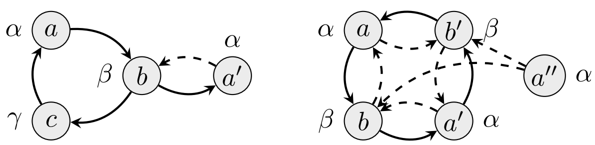

Figure 3: CAFs used as counter examples for I-maximality (cf. Theorem 6). Dashed arrows are edges in $wfp(F)$.

5 I-Maximality

One of the advantages of wfCAFs over general CAFs is that they preserve I-maximality under most maximization-based semantics (cf. Table 1), which leads to more intuitive behavior of these semantics when considering extensions on the claim-level. We now investigate whether these advantages are preserved when preferences are introduced. Analogously to Definition 5, σ_p^i is I-maximal for a class $\mathcal{F}$ of PCAFs if, for all F in $\mathcal{F}$ and all $S, T \in \sigma_p^i(F)$, $S \subseteq T$ implies $S = T$.

From known properties of wfCAFs (cf. Table 1) it follows directly that $naive_p^i$ is not I-maximal for PCAFs. It remains to investigate I-maximality of prf_p^i, stb_p^i, sem_p^i, and stg_p^i. As it turns out, Reduction 3 manages to preserve I-maximality in most cases. Interestingly, the other three reductions lose I-maximality for *all* semantics.

Theorem 6. *The results in Table 3 hold, even when considering only PCAFs with transitive preferences.*

Intuitively, the above result can be explained by the fact that Reduction 3 is the most conservative of the reductions, not adding new attacks and preserving conflict-freeness (i.e., given a PCAF F, a set of arguments E is conflict-free in the underlying CAF of F iff E is conflict-free in $\mathcal{R}_3(F)$). Reductions 2 and 4 preserve conflict-freeness too, but they may introduce new attacks in contrast to Reduction 3. It is easy to see that Reduction 1 does not preserve conflict-freeness. In fact, Reduction 1 has been deemed problematic for exactly this reason when applied to regular AFs (Amgoud and Vesic 2014), although it is still discussed and considered in the literature alongside the other reductions (Kaci et al. 2021).

Let us explore the negative results from Theorem 6 more in-depth. For $i \in \{1, 4\}$, a counter-example is given by the CAF F' from Example 1. For all considered semantics, F' admits the extensions $\{\alpha\}, \{\alpha, \beta\}$ which are in $\subseteq$-relation to each other. It can be checked that $F' \in \mathcal{R}_i\text{-}\mathbf{CAF}_{tr}$.

For $i = 2$, let G be the CAF shown on the right in Figure 3. $G \in \mathcal{R}_2\text{-}\mathbf{CAF}_{tr}$ since $\mathcal{R}_2(G') = G$ for the PCAF G' with attacks $\{(b, a), (b, a'), (b', a), (b', a')\}$ and preferences $a \succ b$ and $a' \succ b'$. Moreover, $stb_c(G) = \{\{\alpha\}, \{\alpha, \beta\}\}$.

For stg_p^3, let H be the CAF shown on the left in Figure 3. Note that $H \in \mathcal{R}_3\text{-}\mathbf{CAF}_{tr}$ and $stg_c^3(H) = \{\{\alpha\}, \{\alpha, \gamma\}, \{\beta\}\}$.

σ	$i=1$	$i \in \{2,4\}$	$i=3$
$cf/adm/naive/stb$	NP-c	in P	in P
com	NP-c	NP-c	in P
$prf/sem/stg$	Σ_2^{P}-c	coNP-c	coNP-c

Table 4: Complexity of $Ver_{\sigma,i}^{PCAF}$.

6 Computational Complexity

In this section, we investigate the impact of preferences on the computational complexity of claim-based reasoning. For each preference reduction $i \in \{1, 2, 3, 4\}$ we define $Cred_{\sigma,i}^{PCAF}$, $Skept_{\sigma,i}^{PCAF}$, and $Ver_{\sigma,i}^{PCAF}$ analogously to $Cred_{\sigma}^{CAF}$, $Skept_{\sigma}^{CAF}$, and Ver_{σ}^{CAF} (cf. Section 2), except that we take a PCAF instead of a CAF as input and appeal to the σ_p^i semantics instead of the σ_c semantics. Membership results for PCAFs can be inferred from results for general CAFs (recall that the preference reductions from PCAFs to CAFs can be done in polynomial time), and hardness results from results for wfCAFs. Thus, the complexity of credulous and skeptical acceptance follows immediately from known results for CAFs and wfCAFs: given $i \in \{1, 2, 3, 4\}$ and $\sigma \in \{cf, adm, com, naive, stb, prf, sem, stg\}$, the problems $Cred_{\sigma,i}^{PCAF}$ and $Skept_{\sigma,i}^{PCAF}$ have the same complexity as $Cred_{\sigma}^{wfCAF}$ and $Skept_{\sigma}^{wfCAF}$ respectively (cf. Table 2).

The computational complexity of the verification problem, on the other hand, is one level higher for general CAFs when compared to wfCAFs (cf. Table 2), i.e., the bounds that existing results yield for PCAFs are not tight. We address this open problem and comprehensively analyze $Ver_{\sigma,i}^{PCAF}$ for each of the considered reductions and semantics.

Theorem 7. *The complexity results in Table 4 hold, even when considering only PCAFs with transitive preferences.*

Observe that when using Reduction 1 we obtain the same complexity as for general CAFs, i.e., the benefits of wfCAFs are lost in this case. On the other hand, Reductions 2–4 preserve the lower complexity of wfCAFs for almost all semantics. As we will see, this can be explained by the fact that these reductions do not remove conflicts between arguments. The only outlier is complete semantics, for which verification remains hard under Reductions 2 and 4 but not Reduction 3. Here, the fact that Reductions 2 and 4 can introduce new attacks leads to an increase in complexity.

We now examine why verification remains easier under Reductions 2–4 in more detail. Given a wfCAF F and a set of claims C, a set of arguments S can be constructed in polynomial time such that S is the unique maximal admissible set in F with claim $cl(S) = C$ (Dvořák and Woltran 2020). Making use of the fact that Reductions 2–4 do not alter conflicts between arguments, we can construct such a maximal set of arguments also for PCAFs: given a PCAF F and set C of claims, we define the set $E_0(C)$ containing all arguments of F with a claim in C; the set $E_1^i(C)$ is obtained from $E_0(C)$ by removing all arguments attacked by $E_0(C)$ in the underlying CAF of F; finally, the set $E_*^i(C)$ is obtained by repeatedly removing all arguments not defended by $E_1^i(C)$ in $\mathcal{R}_i(F)$ until a fixed point is reached.

Definition 11. *Given a PCAF $F = (A, R, cl, \succ)$, a set of claims C, and $i \in \{2, 3, 4\}$, we define*

$$E_0(C) = \{a \in A \mid cl(a) \in C\};$$
$$E_1^i(C) = E_0(C) \setminus E_0(C)^+_{(A,R)};$$
$$E_k^i(C) = \{x \in E_{k-1}^i(C) \mid x \text{ is defended by } E_{k-1}^i(C) \text{ in } \mathcal{R}_i(F)\} \text{ for } k \geq 2;$$
$$E_*^i(C) = E_k^i \text{ for } k \geq 2 \text{ such that } E_k^i(C) = E_{k-1}^i(C).$$

The above definition is based on (Dvořák and Woltran 2020, Definition 5), but with the crucial differences that undefended arguments are (i) computed w.r.t. $\mathcal{R}_i(F)$ and (ii) are iteratively removed until a fixed point is reached.

Lemma 8. *Let F be a PCAF, C a set of claims, and $i \in \{2, 3, 4\}$. The following holds:*

- *$C \in cf_p^i(F)$ iff $cl(E_1^i(C)) = C$. Moreover, if $C \in cf_p^i(F)$ then $E_1^i(C)$ is the unique maximal conflict-free set S in $\mathcal{R}_i(F)$ such that $cl(S) = C$;*
- *$C \in adm_p^i(F)$ iff $cl(E_*^i(C)) = C$. If $C \in adm_p^i(F)$ then $E_*^i(C)$ is the unique maximal admissible set S in $\mathcal{R}_i(F)$ such that $cl(S) = C$.*

By computing the maximal conflict-free (resp. admissible) extensions $E_1^i(C)$ (resp. $E_*^i(C)$) for a claim set C, verification becomes easier for most semantics. For instance, to decide whether $C \in stb_p^i(F)$ we first check if $C \in adm_p^i(F)$. If not, $C \notin stb_p^i(F)$. If yes, then $cl(E_*^i(C)) = C$ (cf. Lemma 8). We check (in P) if $E_*^i(C)$ is stable in $\mathcal{R}_i(F)$. If yes, we are done. If no, there is an argument x that is neither in $E_*^i(C)$ nor attacked by $E_*^i(C)$ in $\mathcal{R}_i(F)$. Moreover, there can be no other stable set S with $cl(S) = C$ in $\mathcal{R}_i(F)$ since $S \subseteq E_*^i(C)$ (cf. Lemma 8).

The lower complexity of the verification problem is crucial for enumerating extensions. In particular, the improved enumeration algorithm for wfCAFs (Dvořák and Woltran 2020) is based on the polynomial time verification of claim-sets and thus extends to PCAFs under Reductions 2–4. This further implies that deciding the main decision problems is tractable if the number of claims is bounded by a constant k, i.e., these problems are fixed parameter tractable (FPT).

Theorem 9. *For $\sigma \in \{cf, adm, naive, stb, prf, sem, stg\}$, $i \in \{2, 3, 4\}$, and for $\sigma = com$ in case $i = 3$, there is a polynomial $poly(\cdot)$ such that $Cred^{PCAF}_{\sigma,i}$, $Skept^{PCAF}_{\sigma,i}$, and $Ver^{PCAF}_{\sigma,i}$ can be solved in time $\mathcal{O}(4^k \cdot poly(n))$ for PCAFs $(A, R, cl, \succ)$ with $|cl(A)| \leq k$.*

7 Conclusion

Many approaches to argumentation (i) assume that arguments with the same claims attack the same arguments and (ii) take preferences into account. Investigations on CAFs so far only consider (i), showing that wfCAFs have several desired properties. In this paper, we tackle (ii) and analyze whether these properties still hold when preferences are taken into account. To this end, we introduced Preference-based CAFs (PCAFs) and investigated the impact of the four commonly used preference reductions on PCAFs.

We examined and characterized resulting CAF-classes, yielding insights into the expressiveness of argumentation formalisms that can be instantiated as CAFs and allow for preference incorporation. Furthermore, we investigated the properties of I-maximality and computational complexity for PCAFs. Preserving I-maximality is desirable since it implies intuitive behavior of maximization-based semantics, while the complexity of the verification problem is crucial for the enumeration of claim-extensions. Insights in terms of both semantical and computational properties provide necessary foundations towards a practical realization of this particular argumentation paradigm (we refer to, e.g., (Baumeister et al. 2021; Fazzinga, Flesca, and Furfaro 2020), for a similar research endeavor in terms of incomplete AFs).

Our results show that (i) Reduction 3 exhibits the same properties as wfCAFs regarding computational complexity and also preserves I-maximality for most semantics; (ii) Reductions 2 and 4 retain the advantages of wfCAFs regarding complexity for all but complete semantics, but do not preserve I-maximality; (iii) under Reduction 1, neither complexity properties nor I-maximality are preserved. The above results hold even if we restrict ourselves to transitive preferences. It is worth noting that Reduction 3 behaves favorably on regular AFs as well, fulfilling many principles for preference-based semantics laid out by Kaci et al. (2021).

In this work, we dealt with preferences via preference reductions that modify the attack relation. Another approach is to lift orderings over arguments to sets of arguments and select extensions in this way (Brewka, Truszczynski, and Woltran 2010; Amgoud and Vesic 2014; Kaci, van der Torre, and Villata 2018; Alfano et al. 2022). These two paradigms interpret the meaning of preferences between arguments differently: using reductions, $x \succ y$ expresses that x is stronger than y, while in the second approach $x \succ y$ expresses that it is preferred to have outcomes with x rather than with y. Interestingly, under Reduction 3, the admissible/complete/stable extensions of a preference-based AF are also extensions in the underlying AF (Kaci et al. 2021). Thus, Reduction 3 selects the 'best' extensions from the underlying AF in these cases. A similar dichotomy concerning preference handling can be observed in related areas such as logic programs, where preferences are incorporated either on the syntactic level (Delgrande, Schaub, and Tompits 2003) or by ranking the outcome (Sakama and Inoue 2000). An interesting avenue for future work is to examine the effect of preference liftings on (well-formed) CAFs.

Another possibility for future work is to extend our studies to alternative semantics for CAFs (Dvořák, Rapberger, and Woltran 2020; Dvořák et al. 2021), where subset-maximization is handled on the claim-level instead of on the argument-level. Experimental evaluation of our results may also be interesting. Lastly, one can lower the level of abstraction used here, e.g., by incorporating more structure into arguments, by allowing arguments to act in support of other arguments as is done in bipolar AFs (Amgoud et al. 2008), or by preserving more information about the claims of arguments. Regarding the latter point, recent research (Wakaki 2020) has shown that formalisms which permit strong negation require careful examination with regards to consistency.

Acknowledgments

We thank the anonymous reviewers for their valuable and constructive feedback. This work was funded by the Austrian Science Fund (FWF) under grants Y698 and W1255-N23 and by the Vienna Science and Technology Fund (WWTF) under grant ICT19-065.

References

Alfano, G.; Greco, S.; Parisi, F.; and Trubitsyna, I. 2022. On Preferences and Priority Rules in Abstract Argumentation. In *Proc. IJCAI'22*, 2517–2524. IJCAI Organization.

Amgoud, L.; and Cayrol, C. 1998. On the Acceptability of Arguments in Preference-based Argumentation. In *Proc. UAI'98*, 1–7. Morgan Kaufmann.

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

Amgoud, L.; Cayrol, C.; Lagasquie, M.-C.; and Livet, P. 2008. On Bipolarity in Argumentation Frameworks. *Int. J. Intelligent Systems*, 23: 1–32.

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

Atkinson, K.; Baroni, P.; Giacomin, M.; Hunter, A.; Prakken, H.; Reed, C.; Simari, G. R.; Thimm, M.; and Villata, S. 2017. Towards Artificial Argumentation. *AI Magazine*, 38(3): 25–36.

Atkinson, K.; and Bench-Capon, T. J. M. 2021. Value-based Argumentation. In *Handbook of Formal Argumentation, Volume 2*, 299–354. College Publications. Also appears in IfCoLog Journal of Logics and their Applications 8(6):1543–1588.

Baroni, P.; Caminada, M.; and Giacomin, M. 2018. Abstract Argumentation Frameworks and Their Semantics. In *Handbook of Formal Argumentation*, chapter 4, 159–236. College Publications.

Baroni, P.; and Giacomin, M. 2007. On principle-based evaluation of extension-based argumentation semantics. *Artif. Intell.*, 171(10-15): 675–700.

Baroni, P.; and Riveret, R. 2019. Enhancing Statement Evaluation in Argumentation via Multi-labelling Systems. *J. Artif. Intell. Res.*, 66: 793–860.

Baumeister, D.; Järvisalo, M.; Neugebauer, D.; Niskanen, A.; and Rothe, J. 2021. Acceptance in incomplete argumentation frameworks. *Artif. Intell.*, 295: 103470.

Bistarelli, S.; and Santini, F. 2021. Weighted Argumentation. In *Handbook of Formal Argumentation, Volume 2*, 355–395. College Publications. Also appears in IfCoLog Journal of Logics and their Applications 8(6):1589–1622.

Brewka, G.; Truszczynski, M.; and Woltran, S. 2010. Representing Preferences Among Sets. In *Proc. AAAI'10*, 273–278. AAAI Press.

Caminada, M.; Sá, S.; Alcântara, J.; and Dvořák, W. 2015. On the equivalence between logic programming semantics and argumentation semantics. *Int. J. Approx. Reasoning*, 58: 87–111.

Cyras, K.; and Toni, F. 2016. ABA+: Assumption-Based Argumentation with Preferences. In *Proc.KR'16*, 553–556. AAAI Press.

Delgrande, J. P.; Schaub, T.; and Tompits, H. 2003. A framework for compiling preferences in logic programs. *Theory and Practice of Logic Programming*, 3(2): 129–187.

Dung, P. M. 1995. On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games. *Artif. Intell.*, 77(2): 321–358.

Dung, P. M.; Kowalski, R. A.; and Toni, F. 2009. Assumption-Based Argumentation. In *Argumentation in Artificial Intelligence*, 199–218. Springer.

Dunne, P. E.; Dvořák, W.; Linsbichler, T.; and Woltran, S. 2015. Characteristics of multiple viewpoints in abstract argumentation. *Artif. Intell.*, 228: 153–178.

Dvořák, W.; Greßler, A.; Rapberger, A.; and Woltran, S. 2021. The Complexity Landscape of Claim-Augmented Argumentation Frameworks. In *Proc. AAAI'21*, 6296–6303. AAAI Press.

Dvořák, W.; Rapberger, A.; and Woltran, S. 2020. Argumentation Semantics under a Claim-centric View: Properties, Expressiveness and Relation to SETAFs. In *Proc. KR'20*, 341–350. IJCAI.org.

Dvořák, W.; and Woltran, S. 2020. Complexity of abstract argumentation under a claim-centric view. *Artif. Intell.*, 285: 103290.

Fazzinga, B.; Flesca, S.; and Furfaro, F. 2020. Revisiting the Notion of Extension over Incomplete Abstract Argumentation Frameworks. In *Proc. IJCAI'20*, 1712–1718.

Horty, J. F. 2002. Skepticism and floating conclusions. *Artif. Intell.*, 135(1-2): 55–72.

Kaci, S.; van der Torre, L. W. N.; Vesic, S.; and Villata, S. 2021. Preference in Abstract Argumentation. In *Handbook of Formal Argumentation, Volume 2*, 211–248. College Publications.

Kaci, S.; van der Torre, L. W. N.; and Villata, S. 2018. Preference in Abstract Argumentation. In *Proc. COMMA'18*, volume 305 of *FAIA*, 405–412. IOS Press.

Li, H.; Oren, N.; and Norman, T. J. 2011. Probabilistic Argumentation Frameworks. In *Proc. TAFA'11*, volume 7132 of *LNCS*, 1–16. Springer.

Modgil, S. 2009. Reasoning about Preferences in Argumentation Frameworks. *Artif. Intell.*, 173(9-10): 901–934.

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

Modgil, S.; and Prakken, H. 2018. Abstract rule-based argumentation. In *Handbook of Formal Argumentation*, chapter 6, 287–364. College Publications.

Prakken, H.; and Vreeswijk, G. A. 2002. Logics for defeasible argumentation. In *Handbook of Philosophical Logic*, volume 4. Springer.

Rapberger, A. 2020. Defining Argumentation Semantics under a Claim-centric View. In *Proc. STAIRS'20*, volume 2655 of *CEUR Workshop Proceedings*. CEUR-WS.org.

Rocha, V. H. N.; and Cozman, F. G. 2022. A Credal Least Undefined Stable Semantics for Probabilistic Logic Programs and Probabilistic Argumentation. In *Proc. KR'22*, 309–319. IJCAI Organization.

Sakama, C.; and Inoue, K. 2000. Prioritized logic programming and its application to commonsense reasoning. *Artif. Intell.*, 123(1-2): 185–222.

van der Torre, L.; and Vesic, S. 2017. The Principle-Based Approach to Abstract Argumentation Semantics. *IfCoLog Journal of Logic and its Applications*, 4(8): 2735–2778.

Wakaki, T. 2020. Consistency in Assumption-Based Argumentation. In *Proc.COMMA'20*, volume 326 of *FAIA*, 371–382. IOS Press.