

Preference in Abstract Argumentation

Souhila KACI^a, Leendert VAN DER TORRE^b and Serena VILLATA^c

^aLIRMM, France

^bComputer Science and Communications, University of Luxembourg, Luxembourg

^cUniversité Côte d'Azur, CNRS, Inria, I3S, France

Abstract. Consider an argument A that is attacked by an argument B , while A is preferred to B . Existing approaches will either ignore the attack or reverse it. In this paper we introduce a new reduction of preference and attack to defeat, based on the idea that in such a case, instead of ignoring the attack, the preference is ignored. We compare this new reduction with the two existing ones using a principle-based approach, for the four Dung semantics. The principle-based or axiomatic approach is a methodology to choose an argumentation semantics for a particular application, and to guide the search for new argumentation semantics. For this analysis, we also introduce a fourth reduction, and a semantics for preference-based argumentation based on extension selection. Our classification of twenty alternatives for preference-based abstract argumentation semantics using six principles suggests that our new reduction has some advantages over the existing ones, in the sense that if the set of preferences increases, the sets of accepted arguments increase as well.

Keywords. Formal argumentation, abstract argumentation, preference, axiomatic approach, principle-based approach

1. Introduction

Preferences are used in abstract argumentation to represent the comparative strength of arguments. Roughly, there are two positions on preference-based argumentation (PAF). One position handles preference at the abstract level of Dung's argumentation framework (AF), assuming a preference order between arguments. Another position is skeptical about this approach, on the grounds that the assumptions made at the abstract level may not hold when considering preferences at the structured level [12]. As arguments are treated in an absolutely abstract way, the proponents of the second position claim, it is no wonder that strange things may happen.

This paper takes the first point of view, and to address the above criticism, it puts some order in the study of how PAFs behave by adopting a principle-based approach. Amgoud and Vesic [3] propose to study preference-based argumentation using a principle-based approach, which was developed more generally for extension-based abstract argumentation semantics by Baroni and Giacomin [5,14]. In particular, Amgoud and Vesic study so-called *Conflict-freeness* and *Generalisation*. We extend their study with a new principle, called *Extension Growth*: if we add preferences, then we can infer more. There are various reasons why the intersection of the extensions grows: either there are less extensions, or the extensions become larger, or both. We therefore study also principles that consider only one of these two aspects, such as *Extension Selection*.

2. Preference-based Argumentation

In Dung's approach [9] the acceptance of an argument depends only on the defeat relation among arguments and the chosen argumentation semantics. The outcome of an argumentation framework is a set of sets of arguments, called *extensions* and denoted by $E(\mathcal{A}, \text{Def})$, that are robust against defeats. Dung distinguishes several definitions of extension (e.g. complete, grounded, preferred, stable), each corresponding to an acceptability semantics that formally rules the argument evaluation process [9].

Definition 1 (Argumentation framework [9]) An argumentation framework (AF) is a tuple $\langle \mathcal{A}, \text{Def} \rangle$ where \mathcal{A} is a set of arguments and $\text{Def} \subseteq \mathcal{A} \times \mathcal{A}$ is a binary defeat relation. For $a, b \in \mathcal{A}$, $(a, b) \in \text{Def}$ stands for a defeats b .

The defeat relation was called "attack" relation by Dung [9], but in preference-based argumentation this name is used for another relation.

Definition 2 (Preference-based argumentation framework [1]) A preference-based argumentation framework (PAF) is a 3-tuple $\langle \mathcal{A}, \text{Att}, \succ \rangle$ where \mathcal{A} is a set of arguments, Att is a binary attack relation $\subseteq \mathcal{A} \times \mathcal{A}$ and \succ is a partial order (irreflexive and transitive) over \mathcal{A} , called preference relation. For $a, b \in \mathcal{A}$, $(a, b) \in \text{Att}$ stands for a attacks b .

To compute the extensions of a preference-based argumentation framework $\langle \mathcal{A}, \text{Att}, \succ \rangle$ the latter can be reduced to a Dung's AF $\langle \mathcal{A}, \text{Def} \rangle$. We say that $\langle \mathcal{A}, \text{Att}, \succ \rangle$ represents $\langle \mathcal{A}, \text{Def} \rangle$. The extensions of a preference-based argumentation framework, denoted by $\mathcal{E}(\mathcal{A}, \text{Att}, \succ)$, are simply the extensions of the argumentation framework it represents. Definition 3 provides two ways proposed to reduce a PAF into Dung's AF. Reduction 1 has been commonly used in almost all approaches to preference-based argumentation. An attack succeeds only when the attacked argument is not preferred to the attacker. This reduction has been criticised by Amgoud and Vesic [3] as it may lead to non conflict-free extensions. The problem occurs when there is an attack from an argument to a preferred argument. This attack is called *critical* by Amgoud and Vesic [3]. Amgoud and Vesic [3] have proposed to repair the argumentation framework. They extended Reduction 1 by enforcing a defeat from an argument to another when the former is preferred but attacked by the latter. This is Reduction 2.

Definition 3 (Existing reductions of PAF to AF) Let $\langle \mathcal{A}, \text{Att}, \succ \rangle$ be a PAF and $\langle \mathcal{A}, \text{Def} \rangle$ be the AF it represents.

- Reduction 1 [1]: $\forall a, b \in \mathcal{A} : (a, b) \in \text{Def} \text{ iff } (a, b) \in \text{Att}, b \not\succeq a$.
- Reduction 2 [3]: $\forall a, b \in \mathcal{A} : (a, b) \in \text{Def} \text{ iff } ((a, b) \in \text{Att}, b \not\succeq a) \text{ or } ((b, a) \in \text{Att}, (a, b) \notin \text{Att}, a \succ b)$.

Example 1 (Reduction 1 and 2) Let $\langle \mathcal{A}, \text{Att}, \succ \rangle$ with $\mathcal{A} = \{a, b\}$, $\text{Att} = \{(a, b)\}$ and $b \succ a$. With Reduction 1 we have $\text{Def} = \emptyset$. Both a and b are accepted using any semantics although they are conflicting w.r.t. Att . With reduction 2 we have $(b, a) \in \text{Def}$. b is accepted.

Reduction 2 is based on an implicit strong constraint that an argument never succeeds to attack a preferred argument. This view gives a *power* to preferred arguments. It is

easy to construct examples where this is counterintuitive. For example, consider a parent who refuses that his child watches TV in the evening during the week because he has courses. However, the child says that his courses have been cancelled. Then to maintain the refusal to watch TV the parent should provide another argument attacking his child’s argument.

It is well known from other areas of knowledge representation, that it is not too hard to find a counterexample for any kind of representation. In other words, sometimes we have to rephrase examples such that they give the desired conclusions. However, the example suggests that there are straightforward alternatives to Reduction 1 or 2. Moreover, Reduction 3 can be criticized arguing that it would not be natural to make successful an attack from a less preferred argument. Reduction 4 below mixes Reduction 2 and 3.


Definition 4 (New reductions from PAF to AF) Let $\langle \mathcal{A}, \text{Att}, \succ \rangle$ be a PAF and $\langle \mathcal{A}, \text{Def} \rangle$ be the AF it represents.

- Reduction 3 $\forall a, b \in \mathcal{A} : (a, b) \in \text{Def} \text{ iff } ((a, b) \in \text{Att}, b \not\succeq a) \text{ or } ((a, b) \in \text{Att}, (b, a) \notin \text{Att}).$
- Reduction 4 $\forall a, b \in \mathcal{A} : (a, b) \in \text{Def} \text{ iff } ((a, b) \in \text{Att}, b \not\succeq a) \text{ or } ((b, a) \in \text{Att}, (a, b) \notin \text{Att}, a \succ b) \text{ or } ((a, b) \in \text{Att}, (b, a) \notin \text{Att}).$

Example 2 (Example 1 continued – Reduction 3 and 4) With Reduction 3 we have $(a, b) \in \text{Def}$. a is accepted. With Reduction 4 we have $(a, b) \in \text{Def}$ and $(b, a) \in \text{Def}$. The grounded extension is empty, complete extensions are $\{\}, \{a\}, \{b\}$, and there are two preferred/stable extensions $\{a\}$ and $\{b\}$.

The following example further illustrates the four reductions.

Example 3 Assume the PAF visualised below, where we added an argument c , and b attacks c , c attacks a . For reduction 1, 2, 3 and 4 respectively, as visualised below, PAF semantics give $\{\{a, b\}\}, \{\{b\}\}$, no extensions for stable and $\{\{\}\}$ for the rest, and $\{\{b\}, \{\}\}$ for complete, the first for preferred and stable, the latter for grounded. Note that when no preference is assumed among arguments, then no argument is accepted, i.e., the grounded extension is empty. Now, when a preference of b over a is added, one would maybe expect to get b accepted, possibly c could be accepted but there does not seem to be any reason to accept a . Nevertheless, Reduction 1 gives the counterintuitive result to accept both a and b .



3. Extension selection based on preferences

To illustrate our principles we introduce a PAF semantics that is not based on a reduction, but on extension selection. Assume a set of extensions, and a preference relation

over arguments. To get the best extensions, we need to lift the preference relation over arguments to a preference relation over sets of arguments, or extensions. It is well known that there are various ways to make this lifting more precise, see for example the work of Amgoud and Vesic [3]. Here we use the following lifting: if argument a is preferred to argument b , then all extensions containing argument a but not b are either better than all extensions containing b but not a , or the two extensions are incomparable.

Definition 5 Let $E(\mathcal{A}, \text{Def})$ be a set of extensions according to a Dung semantics, and let \succ is an order (irreflexive and transitive) over \mathcal{A} , called preference relation. $E \subseteq \mathcal{A}$ is at least as good as $E' \subseteq \mathcal{A}$ if $\forall a, b \in \mathcal{A}$: we do not have that $a \succ b$ and $a \in E' \setminus E$ and $b \in E \setminus E'$. E is better than E' iff E at least as good as E' and E' not at least as good as E . E is best if there is no E' that is better than E .


We can define a semantics of preference-based argumentation by a two-step process. First we select the extensions of the framework according to a regular abstract semantics. Then we use the preference relation to select the best extensions among them. We refer to the following construction as PAF semantics 5. We identify $\text{Def}=\text{Att}$ in Dung’s AF and define $\mathcal{E}(\mathcal{A}, \text{Att}, \succ)$ as a selection of the best extensions of $E(\mathcal{A}, \text{Att})$.

Definition 6 The extensions $\mathcal{E}(\mathcal{A}, \text{Att}, \succ)$ are the best extensions among $E(\mathcal{A}, \text{Att})$ based on \succ .

In preference elicitation, the preferences are extracted from a user step by step. Assume now that for each step where a user is querying for a preference, we consider the arguments the user accepts, given the knowledge about the preference relation thus far. In such a setting, it would be quite useful if the set of accepted arguments is increasing monotonically. Moreover, the other way around, we can even consider scenarios where the interest in the extension guides the order in which the preferences are elicited.

It is important to note that this construction is very different from the reduction-based semantics, and thus leads very different extensions. In Example 4, there is only one extension $\{a\}$, and thus there is only one preferred extension $\{a\}$. In Example 5 there is either no extension (for stable) or one extension $\{\}$ (all other semantics). Thus, in both cases the extension selection-based semantics coincides with reduction 3.

Example 4 Consider the framework below under preferred or stable semantics. The left PAF prefers extension $\{a, c\}$ over extension $\{b, d\}$, but if we add $b \succ a$ then both $\{a, c\}$ and $\{b, d\}$ are best extensions.



This is just one of the many possibilities, and its simplicity has some drawbacks. For example, if the empty set is an extension, e.g. in complete semantics, then it is always a best extension. Just like in the case of the reductions, this illustrates that more alternatives can be defined, and a principle-based approach is needed to choose among the alternatives. We leave a further study of this for future research.

4. Principle-based analysis

In this section, we investigate the definition of some general criteria for evaluating PAFs. Amgoud and Vesic [3] introduce two principles, called conflict-freeness and generalisation. We also use conflict-freeness for our principle-based analysis, but in addition we define principles based on extension growth and selection. Principle 1 is a version of the conflict-freeness principle for preference-based argumentation introduced by Amgoud and Vesic [3]. If a attacks b then there is no extension containing both a and b .

- Principle 1 (P1): Conflict-freeness

$$\text{If } (a, b) \in \text{Att} \text{ then } \nexists E \in \mathcal{E}(\mathcal{A}, \text{Att}, \succ) \mid \{a, b\} \subseteq E.$$

Principle 2 and 3 are inspired by the idea that preferences are used to select the best among all extensions. The rationale behind Principle 4 is that increasing the number of expressed preferences can have the effect that the number of extensions decreases, or that the extension grows. The same holds for Principle 5. Principle 6, instead, states that increasing the number of preferences does not increase the number of extensions.

- Principle 2 (P2): Preference selects extensions 1

$$\mathcal{E}(\mathcal{A}, \text{Att}, \succ \cup \succ') \subseteq \mathcal{E}(\mathcal{A}, \text{Att}, \succ).$$

- Principle 3 (P3): Preference selects extensions 2

$$\mathcal{E}(\mathcal{A}, \text{Att}, \succ) \subseteq \mathcal{E}(\mathcal{A}, \text{Att}, \{ \}).$$

- Principle 4 (P4): Extension refinement

$$\forall E \in \mathcal{E}(\mathcal{A}, \text{Att}, \succ \cup \succ'), \exists E' \in \mathcal{E}(\mathcal{A}, \text{Att}, \succ) \mid E' \subseteq E.$$

- Principle 5 (P5): Extension growth

$$\bigcap \mathcal{E}(\mathcal{A}, \text{Att}, \succ) \subseteq \bigcap \mathcal{E}(\mathcal{A}, \text{Att}, \succ \cup \succ').$$

- Principle 6 (P6): Number of extensions

$$|\mathcal{E}(\mathcal{A}, \text{Att}, \succ \cup \succ')| \leq |\mathcal{E}(\mathcal{A}, \text{Att}, \succ)|.$$

P2 implies P3 and P6 but not vice versa. Moreover, P4 implies P5 but not vice versa. The reader may argue that the way P4 and P5 are formulated resembles the skeptical relations between semantics introduced by Baroni and Giacomin [6]. Investigating how to express principles on the basis of these skeptical relations is left for future work.


Table 1 reports satisfiability of Reductions with respect to the above defined principles, focusing on standard Dung's semantics. The question mark refers to an open problem. If a letter is missing, it means that there is a counterexample. For example, if we consider *Conflict-freeness* (P1) and *Extension Growth* (P5) as required properties, then these results indicate that we either should select R3 or PAF Semantics 5, and we should adopt either grounded or complete semantics.

Red.	P1	P2	P3	P4	P5	P6
R1	×	×	×	GC	GC	G
R2	CGPS	×	×	×	×	G
R3	CGPS	CS	CS	CGS	CGS	CGS P?
R4	CGPS	×	×	×	×	G
S5	CGPS	G	CGPS	CG	CG	G

Table 1. Comparison among the resolutions and the proposed principles. We refer to Dung's semantics as follows: Complete (C), Grounded (G), Preferred (P), Stable (S). When a principle is never satisfied by a certain reduction for all semantics, we use the × symbol. P1 refers to Principle 1, the same holds for the others.


Example 3 is a counterexample for P1 for Reduction 1 for all semantics. The positive results for P1 rely on the notion of *conflict-freeness*. Examples 1-2 are counterexamples for P2 and P3 for Reduction 1 for all semantics, P2 and P5 for Reduction 2 and 4 for all semantics. Example 3 is a counterexample for P2 and P6 for Reduction 1, 2, 4 for stable semantics. The positive results for Reduction 1 and 3 follows from the fact that attack relation is shrinking. Concerning P2 and P3 for Reductions 1-4, the left example in Example 5 is a counterexample for grounded semantics. The right example in Example 5 is a counterexample to P2 and P5 for Reduction 1-4 for preferred semantics.

Example 5 Assume the PAF on the left. Reductions 1-4 coincide. Without preference the grounded extension is $\{\}$, with preference $\{a\}$. Assume the PAF on the right. All reductions coincide. Without preference the unique preferred or stable extension is $\{a\}$, with preference the preferred extension is $\{\}$, and the stable extension does not exist.




The following example is a counterexample to P2 and P3 for Reduction 1 and 2. It is also a counterexample to P2 and P3 for Reduction 4 for all semantics except grounded.

Example 6 Assume the PAF below. Without preference the complete extensions are $\{\}$ and $\{a\}$, the former is the grounded extension and the latter is the preferred or stable extension. For R1, all semantics coincide and the extension is $\{a,b\}$. For R2, all semantics coincide and the extension is $\{b\}$. For R3, the extensions are the same as for the framework without preference. For R4, the complete extensions are $\{\}$, $\{a\}$ and $\{b\}$, the first is the grounded extension, and the latter two are the preferred or stable extensions.



Example 4 shows that PAF Semantics 5 does not satisfy P2 for complete, preferred and stable, and Example 3 shows the same for P6. Moreover, it does not satisfy P4 and P5 for preferred and stable. The positive results for PAF Semantics 5 and P3 follow directly from the definitions. Example 7 is a counterexample to P6 for Reduction 1, 2 and 4, for all semantics except grounded.

Example 7 Assume the PAF below. Without preference the complete extensions are $\{\}$ and $\{a,d\}$, the former is the grounded extension and the latter is the preferred or stable extension. For R1, all semantics coincide, and the extensions are $\{a,b,e\}$, $\{a,b,f\}$ and $\{a,b,g\}$. For R2, all semantics coincide and the extensions are $\{b,e\}$, $\{b,f\}$ and $\{b,g\}$. For R4, the complete extensions are $\{\}$, $\{a,d\}$, $\{b,e\}$, $\{b,f\}$ and $\{b,g\}$.



5. Related work

We distinguish between different proposals to integrate preferences in argumentation. Amgoud *et al.* [2] construct arguments from a set of classical propositional formulas pervaded with priorities, allowing to associate a strength to each argument. The preference relation is defined on the strength of the arguments. In valued-based argumentation [7], arguments are associated to values. The preference relation over arguments is derived from the preference relation over values. The more important the value, the more preferred the argument is. Modgil [11] derives the preference relation from arguments supporting preferences between arguments. The common point of these preference-based argumentation frameworks is that they all rely on Reduction 1, thus violating conflict-freeness of extensions with respect to the attack relation. Amgoud and Vesic [3] have proposed a new PAF which satisfies two requirements: conflict-freeness of extensions with respect to the attack relation (this is Principle 1 in our analysis), and, in the absence of critical attacks, the extensions of a preference-based argumentation framework coincide with the extensions of Dung's argumentation framework. Moreover the PAF handles critical attacks following Reduction 2. The second requirement can be simply reformulated as if the preference relation is empty, meaning that the extensions of a PAF and Dung's framework coincide. This is a special case of Principle 3 in our analysis. To compute extensions, Amgoud and Vesic [3] define a preference relation over the powerset of arguments using the so-called democratic and elitist relations. Maximal conflict-free subsets are extensions of the PAF (grounded, preferred and stable extensions). Our Definition 5 is incomparable with democratic and elitist relations. Definition 5 applies on the PAF with ignoring the associated preference relation while [3] applies democratic and elitist relations on the PAF with Reduction 2.

6. Conclusion

In this paper, we proposed an axiomatic approach to preference-based semantics. We considered four reductions to move from PAFs to a Dung-like abstract argumentation on which standard semantics can be applied to compute the set of accepted arguments, and we proposed a set of six principles we used to study the considered reductions.

The results of this paper give rise to many new research questions. Many more principles can be defined in our framework (e.g., following [13]), and used in the analysis. In particular, it is striking that many principles have a dynamic flavor, and we conjecture that many approaches to dynamics of argumentation [8] can be used as a source for principles. We are in particular interested in principles that distinguish the various

PAF semantics. We believe that there are not many new reductions to be found, but more PAF semantics can be defined not based on reductions. We observe that the resolution-based family of abstract argumentation semantics [4] seems also related to Reduction 3 introduced in this paper, as well as the results of Kaci et al. [10].

Acknowledgements

We thank the anonymous reviewers for valuable comments. Leendert van der Torre and Serena Villata have received funding from the European Union's H2020 research and innovation programme under the Marie Curie grant agreement No. 690974 for the project MIREL: MIning and REasoning with Legal texts.

References

- [1] Leila Amgoud and Claudette Cayrol. Inferring from inconsistency in preference-based argumentation frameworks. *International Journal of Approximate Reasoning*, 29(2):125–169, 2002.
- [2] Leila Amgoud, Claudette Cayrol, and Daniel LeBerre. Comparing arguments using preference orderings for argument-based reasoning. In *ICTAI'96*, pages 400–403, 1996.
- [3] Leila Amgoud and Srdjan Vesic. Rich preference-based argumentation frameworks. *Int. J. Approx. Reasoning*, 55(2):585–606, 2014.
- [4] Pietro Baroni, Paul E. Dunne, and Massimiliano Giacomin. On the resolution-based family of abstract argumentation semantics and its grounded instance. *Artif. Intell.*, 175(3-4):791–813, 2011.
- [5] Pietro Baroni and Massimiliano Giacomin. On principle-based evaluation of extension-based argumentation semantics. *Artif. Intell.*, 171(10-15):675–700, 2007.
- [6] Pietro Baroni and Massimiliano Giacomin. Skepticism relations for comparing argumentation semantics. *Int. J. Approx. Reasoning*, 50(6):854–866, 2009.
- [7] Trevor J.M. Bench-Capon. Persuasion in practical argument using value-based argumentation frameworks. *Journal of Logic and Computation*, 13(3):429–448, 2003.
- [8] Richard Booth, Souhila Kaci, Tjitze Rienstra, and Leendert W. N. van der Torre. A logical theory about dynamics in abstract argumentation. In *SUM'13*, pages 148–161, 2013.
- [9] Phan M. Dung. On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77:321–357, 1995.
- [10] Souhila Kaci, Leendert W. N. van der Torre, and Emil Weydert. Acyclic argumentation: Attack = conflict + preference. In *ECAI'06*, pages 725–726, 2006.
- [11] Sanjay Modgil. Reasoning about preferences in argumentation frameworks. *Artificial Intelligence*, 173(9-10):901–934, 2009.
- [12] Sanjay Modgil and Henry Prakken. Resolutions in structured argumentation. In *COMMA'12*, pages 310–321, 2012.
- [13] Tjitze Rienstra, Chiaki Sakama, and Leendert W. N. van der Torre. Persistence and monotony properties of argumentation semantics. In *TAFI'15*, pages 211–225, 2015.
- [14] Leendert van der Torre and Srdjan Vesic. The principle-based approach to abstract argumentation. In Pietro Baroni, Dov Gabbay, Massimiliano Giacomin, and Leendert van der Torre, editors, *The Handbook of Formal Argumentation*. College Publications, 2018.