Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning (KR 2020)
Main Track

Argument, I Choose You! Preferences and Ranking Semantics in Abstract Argumentation

Jean-Guy Mailly[1] , **Julien Rossit**[1]
[1]LIPADE, Université de Paris
{jean-guy.mailly,julien.rossit}@u-paris.fr

Abstract

Preference-based argumentation and ranking semantics are two important research topics in the computational argumentation literature. Surprisingly, no study investigates to what extent preferences over arguments and ranking semantics can interact. This paper fills the gap between the relative priorities that one can express and the evaluation of arguments individual acceptability. More precisely, we propose a natural principle that should be satisfied by a ranking semantics for Preference-based Argumentation Frameworks. We show that although existing semantics do not satisfy this desirable principle, they can be used to define new ranking semantics that exhibit the expected behavior. Finally, we discuss an application of these semantics to the modeling of human reasoning.

1 Introduction

The introduction of preferences in the reasoning process has been an important topic in artificial intelligence (Pigozzi, Tsoukiàs, and Viappiani 2016). In particular, several approaches tackle the question of preferences in abstract argumentation (Amgoud and Cayrol 2002; Amgoud and Vesic 2014; Kaci, van der Torre, and Villata 2018), thus defining generalized versions of the Dung's abstract argumentation framework (AF) (Dung 1995). The semantics of these preference-based AFs (PAFs) are defined following two directions. First, the attack relation can be combined with the preference relation to define what is called a defeat relation. Then, the classical semantics are applied to the defeat graph. The second approach consists in computing extensions classically on the attack graph, then refining these extensions with the preference relation. Surprisingly, there has been no study of preferences in the context of ranking or gradual semantics. These semantics were proposed more recently and have received much attention since then (*e.g.* (Besnard and Hunter 2001; Cayrol and Lagasquie-Schiex 2005; Matt and Toni 2008; Amgoud and Ben-Naim 2013; Bonzon et al. 2016; Grossi and Modgil 2019)).

In this paper, we focus on ranking semantics, *i.e.* semantics that associate argumentation frameworks with rankings allowing the comparison of arguments individual acceptability. We show that existing semantics are not suited to situations where preferences are provided, but we can use them as a base for proposing the first approach that combines ranking semantics and preferences in abstract argumentation. In this new approach, there is a direct correlation between the fact that an argument a is preferred to an argument b, and the fact that a should be more acceptable than b. This principle is directly related to the way people analyze debates when they have preferences about the (sources of) arguments (Krauss 1988; Jamieson and Birdsell 1988), for instance in a political context. This may seem to be a cognitive bias, and indeed we will see that basic principles of ranking semantics are incompatible with our preference handling. However, modeling of human reasoning needs to take into account such biases, especially if we need to predict the behavior and beliefs of a human being (this is, for instance, a crucial question in the context of interactions between human agents and artificial agents (Rosenfeld and Kraus 2018)). Let us recall that, when we consider extension semantics (*i.e.* joint acceptability of arguments), empirical studies show that even the most basic rationality principles of semantics are not satisfied by people. For instance, (Rosenfeld and Kraus 2014) shows that 22% of the test subjects violate conflict-freeness, and 51% violate admissibility. Similarly, the recent study by (Cramer and Guillaume 2019) reports between 24.54% and 37.21% of difference between the arguments accepted by test subjects and the arguments accepted under difference extension semantics. Thus, it makes sense to study ranking semantics that are incompatible with the "classical" principles if we want to consider human beings involved in the process.

The paper is organized as follows. The next section recalls the basic notions of abstract argumentation, especially ranking semantics and existing approaches for handling preferences in abstract argumentation. Then, we propose a principle that describes the behavior expected from a semantics for PAFs, and we show that existing semantics do not satisfy it. Afterwards, we define a family of ranking semantics for PAFs that combine the preference relation and existing ranking semantics in a way that guarantees they satisfy our new principle. Finally, we describe an example of application of our new semantics on a concrete scenario, before concluding the paper by highlighting some future work.

2 Background and Related Work

Let us first recall the basic notions of abstract argumentation:

Definition 1 ((Dung 1995))**.** An Argumentation Framework (AF) is a pair $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$ where $\mathcal{A}$ is a set of arguments and $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$ an attack relation over these arguments.

In the scope of this paper, we only consider classical argumentation frameworks, *i.e.* there are no weights attached to arguments or attacks. Similarly, we only consider attacks between arguments, support relations are put aside for further study.

The initial method for evaluating argument acceptability is based on the notion of extensions, *i.e.* sets of arguments that are jointly acceptable. For an extension semantics σ, $\sigma(\mathcal{F})$ is then a set of extensions. An argument is skeptically accepted if it belongs to each extension, rejected if it belongs to no extension, and credulously accepted otherwise. The set of skeptically accepted (resp. rejected, credulously accepted) arguments is denoted by $\textsc{Sk}_\sigma(\mathcal{F})$ (resp. $\textsc{Rej}_\sigma(\mathcal{F})$, $\textsc{Cr}_\sigma(\mathcal{F})$).

The exact definition of extension semantics is out of the scope of this paper, so we refer the reader to (Dung 1995; Baroni, Caminada, and Giacomin 2018) for an overview.

More recently, an individual view of argument acceptability has been proposed: instead of mapping an AF with a set of extensions, a semantics maps an AF with:

- either an acceptance degree (that can be modeled as a function ν from arguments to real numbers: $\nu(a) \geq \nu(b)$ means that a is at least as acceptable as b);
- or a relative acceptance ranking.

We call the former *a gradual semantics*, and the latter *a ranking semantics*.[1] Since the acceptance degree function of a gradual semantics can be associated with a ranking of arguments, we focus on ranking semantics in this paper.

Definition 2. A ranking semantics σ is a mapping from any AF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$ to a pre-order $\geq_\sigma$ on $\mathcal{A}$.

Classically, $a \simeq_\sigma b$ means ($a \geq_\sigma b$ and $b \geq_\sigma a$), and $a >_\sigma b$ means ($a \geq_\sigma b$ and $b \not\geq_\sigma a$).

The notion of preferences in argumentation has been tackled in different ways in the context of extension-based semantics. A Preference-based Argumentation Framework (PAF) is a triple $\mathcal{P} = \langle \mathcal{A}, \mathcal{R}, \succeq_p \rangle$ where $\mathcal{A}$ and $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$ are the classical parts of an AF, and $\succeq_p \subseteq \mathcal{A} \times \mathcal{A}$ is the preference relation over arguments.

In some existing works, the preference relation is a partial (strict) order (Amgoud and Cayrol 2002; Kaci, van der Torre, and Villata 2018). In this case, if an agent cannot distinguish between two arguments, then they are incomparable with respect to the preference relation, and there is no notion of "equivalent" arguments. Here, we make a different assumption, and define the preference relation as a partial pre-order, similarly to the approach followed in (Pigozzi, Tsoukiàs, and Viappiani 2016):

- the relation is transitive *i.e.* $\forall a, b, c \in \mathcal{A}$, if $a \succeq_p b$ and $b \succeq_p c$, then $a \succeq_p c$;
- the relation is reflexive *i.e.* $\forall a \in \mathcal{A}$, $a \succeq_p a$.

Again, we use $\succ_p$ and $\simeq_p$ to denote respectively the strict and symmetric counterparts of $\succeq_p$, *i.e.* $a \succ_p b$ iff ($a \succeq_p b$ and $b \not\succeq_p a$) and $a \simeq_p b$ iff ($a \succeq_p b$ and $b \succeq_p a$).

[1] Through a misuse of language, we use the word *ranking* for both total and partial pre-orders.

Obviously, if preferences are expressed as a partial order, this is a special case of partial pre-order. So any strict partial preference relation $\succ$ can be modelled in a non-strict relation: $a \succ b$ implies $a \succ_p b$; and a, b being incomparable with respect to $\succ$ implies $a \not\succeq_p b$ and $b \not\succeq_p a$. In that case, there are no a, b s.t. $a \simeq_p b$.

The classical ways of handling preferences consist in defining a *defeat* relation that is the combination of attacks and preferences, and then applying the usual semantics on the defeat graph (Amgoud and Cayrol 2002; Amgoud and Vesic 2014; Kaci, van der Torre, and Villata 2018). The possible defeat relations are:

- $a \rhd_1 b$ iff $(a, b) \in \mathcal{R}$ and $b \not\succ_p a$ (Amgoud and Cayrol 2002);
- $a \rhd_2 b$ iff $a \rhd_1 b$ or $(b, a) \in \mathcal{R}$ and $a \succ_p b$ (Amgoud and Vesic 2014);
- $a \rhd_3 b$ iff $a \rhd_1 b$ or ($(a, b) \in \mathcal{R}$ and $(b, a) \notin \mathcal{R}$) (Kaci, van der Torre, and Villata 2018);
- $a \rhd_4 b$ iff $a \rhd_3 b$ or ($(b, a) \in \mathcal{R}$, $(a, b) \notin \mathcal{R}$ and $a \succ_p b$) (Kaci, van der Torre, and Villata 2018).

When the preference between two arguments is consistent with the direction of attacks (*i.e.* a attacks b, and a is strictly preferred to b), nothing special is made: this attack holds. The difference between these defeat relations concerns their way of handling *critical attacks*, that are situations where a attacks b and b is strictly preferred to a. The first relation ignores these attacks. Notice that in this case, conflicting arguments may be jointly accepted. The other approaches keep conflict-freeness, since

- the attack is reversed by $\rhd_2$;
- the attack is deleted by $\rhd_3$ only if the opposite attack (b, a) also belongs to the initial AF;
- the attack is made symmetric by $\rhd_4$.

The other way to use preferences in PAFs has been proposed in (Amgoud and Vesic 2014; Kaci, van der Torre, and Villata 2018). Here, the semantics are classically applied on the attack graph associated with the PAF (*i.e.* $\langle \mathcal{A}, \mathcal{R} \rangle$). Then, the preference relation over arguments $\succeq_p$ is used to define a relation over extensions, and the best extensions with respect to this new relation are selected.

3 A Principle for Preferences and Ranking Semantics

We propose a new principle that expresses a desirable property of ranking semantics for PAFs. The intuition is that the preference relation should have a direct impact on the semantics: if an argument a is preferred to an argument b, then a should be at least as acceptable as b. Formally, a ranking semantics σ satisfies *Preference Precedence* if, for any PAF $\mathcal{P} = \langle \mathcal{A}, \mathcal{R}, \succeq_p \rangle$, $\forall a, b \in \mathcal{A}$,

(PP) if $a \succ_p b$, then $a >_\sigma b$.

It seems natural to consider that the acceptance of arguments by an agent should reflect its preferences about arguments. Especially, since it corresponds to the reasoning of human beings, autonomous agents must be able to use this kind of reasoning in order to efficiently interact with people.

In the rest of this section, we show counter-examples proving that neither the existing ranking semantics nor the existing PAF semantics satisfy **(PP)**.

3.1 Ranking Semantics

Direct Application of Ranking Semantics From a technical point of view, we can use any existing ranking semantics σ (defined on AFs) on a PAF: for $\mathcal{P} = \langle \mathcal{A}, \mathcal{R}, \succeq_p \rangle$, $\sigma(\mathcal{P}) = \sigma(\langle \mathcal{A}, \mathcal{R} \rangle)$. However, since this does not take into account the preference relation, this direct adaptation of a ranking semantics to PAFs do not satisfy **(PP)**. This is a direct consequence of the fact that the classical principles of ranking semantics are incompatible with **(PP)**. For instance, many ranking semantics (for AFs) satisfy the *Void Precedence* property: for any AF $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$, $\forall a, b \in \mathcal{A}$,

(VP) if $\nexists c \in \mathcal{A}$ such that $(c, a) \in \mathcal{R}$, and $\exists c \in \mathcal{A}$ such that $(c, b) \in \mathcal{R}$, then $a >_\sigma b$.

Intuitively, **(VP)** states that unattacked arguments should be more acceptable than attacked arguments. Although not all semantics satisfy it (see (Bonzon et al. 2017) for a discussion of the non-necessity of **(VP)** for some application), most of them do (Bonzon et al. 2016). We show that such semantics do not satisfy **(PP)**.

Observation 1. *There is no ranking semantics that satisfies both* **(VP)** *and* **(PP)** *for all* $\mathcal{P} = \langle \mathcal{A}, \mathcal{R}, \succeq_p \rangle$.

For a proof of that, let us consider $\mathcal{P} = \langle \{a, b\}, \{(a, b)\}, \succeq_p \rangle)$ with $b \succ_p a$. For any σ that satisfies **(VP)**, $a >_\sigma b$. This violates **(PP)**.

A similar counter-example can be used to show that **(PP)** is incompatible with other principles of ranking semantics. Informally, here are some of them:

(CP) The more an argument is attacked, the weaker it is.

(QP) If some attacker of a is more acceptable than all attackers of b, then a is more acceptable than b.

(CT) If the attackers of a are at least as many and acceptable as those of b, then b is at least as acceptable than a.

(DP) If two argument have the same number of attackers, the one with more defenders is more acceptable.

Applying Ranking Semantics in Defeat Graphs Another way to use existing ranking semantics on PAFs would be to define a defeat graph, and then to apply the semantics on this graph. But again, this method does not satisfy **(PP)**.

Observation 2. *Applying existing ranking semantics on defeat graphs violates* **(PP)***.*

For $\mathcal{P} = \langle \{a, b\}, \{(a, b)\}, \succeq_p \rangle)$ with $b \succ_p a$, the graph based on $\triangleright_1$ ignores the attack, while the graph based on $\triangleright_4$ makes it symmetric. In both cases, any reasonable ranking semantics should consider a and b as equivalently acceptable, thus **(PP)** is violated. Similar counter-examples can be exhibited for $\triangleright_2$ and $\triangleright_3$.

3.2 Extension Semantics for PAFs

Now we check whether existing semantics for PAFs satisfy **(PP)**. We define a way to compare the individual acceptability of arguments based on the set of extensions.

Definition 3. For $\mathcal{P} = \langle \mathcal{A}, \mathcal{R}, \succeq_p \rangle$ and σ an extension semantics for PAFs, the acceptance status of an argument $a \in \mathcal{A}$ is $\text{ACC}(a) = \text{SK}$ (resp. $\text{ACC}(a) = \text{CR}$, $\text{ACC}(a) = \text{REJ}$) if $a \in \text{SK}_\sigma(\mathcal{P})$ (resp. $a \in \text{CR}_\sigma(\mathcal{P})$, $a \in \text{REJ}_\sigma(\mathcal{P})$). By convention, $\text{SK} > \text{CR} > \text{REJ}$.

We define $\geq_\sigma$ over $\mathcal{A}$ by $a >_\sigma b$ if $\text{ACC}(a) > \text{ACC}(b)$, and $a \simeq_\sigma b$ if $\text{ACC}(a) = \text{ACC}(b)$.

PAF Extension Selection We focus here on the PAFs semantics that use the preference relation to refine the extensions of the attack graph (Amgoud and Vesic 2014; Kaci, van der Torre, and Villata 2018).

Observation 3. *Semantics based on extensions refinement and the rankings from Definition 3 do not satisfy* **(PP)***.*

For any such semantics, the attack graph of the PAF $\mathcal{P}_1 = \langle \{a, b, c\}, \{(a, b), (b, c)\}, \succeq_p \rangle$ shown on Figure 1 has a single extension $\varepsilon = \{a, c\}$ for any reasonable semantics, so ε is also the single extension of the PAF. This means that $a \simeq_\sigma c$ and $a, c >_\sigma b$. It is easy to show some preference relation such that this result violates **(PP)** (*e.g.* $b \succ_p c$ and $c \simeq_p a$).

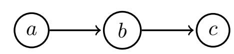

Figure 1: The Attack Graph of the PAF $\mathcal{P}_1$

Defeat-based Semantics We consider now semantics based on defeat graphs and the ranking $\geq_\sigma$ (Definition 3).

Observation 4. *Semantics based on defeat graphs and the rankings from Definition 3 do not satisfy* **(PP)***.*

For $\triangleright_1$, let $\mathcal{P}_2 = \langle \{a, b\}, \{(a, b)\}, \succeq_p \rangle$ be a PAF such that $b \succ_p a$. The defeat relation is empty, so both arguments are skeptically accepted for any reasonable semantics, so **(PP)** is violated. For $\triangleright_2$, we consider again $\mathcal{P}_1$ shown above, with $\succeq_p$ defined by $c \succ_p b \succ_p a$. The defeat graph is shown on Figure 2. For any reasonable semantics the unique extension is $\{a, c\}$, thus a is more acceptable than b, and again **(PP)** is violated. Similar counter-examples can be constructed for proving the result with $\triangleright_3$ and $\triangleright_4$.

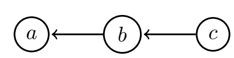

Figure 2: The Defeat Graph of $\mathcal{P}_1$, with $c \succ_p b \succ_p a$

4 Ranking Semantics for PAFs

In this section, we show how to define ranking semantics that satisfy **(PP)**, and we give an example of application to the modeling of human reasoning.

4.1 Preference-Sensitive Semantics for PAFs

The idea is to combine the preference relation $\succeq_p$ and some existing ranking semantics γ to define a new semantics σ. First, if an argument a is strictly preferred to an argument b, then a should be strictly more acceptable than b (with respect to σ). Then, to distinguish between arguments that are equally preferred, we use the γ-ranking to determine which one is more acceptable for σ. Formally:

Definition 4. For $\mathcal{P} = \langle \mathcal{A}, \mathcal{R}, \succeq_p \rangle$ and a ranking semantics γ, we define the semantics σ^p_γ by $\sigma^p_\gamma(\mathcal{P}) = \geq^p_\gamma$ s.t.:
- $a >^p_\gamma b$ iff either $a \succ_p b$, or $a \simeq_p b$ and $a >_\gamma b$;
- $a \simeq^p_\gamma b$ iff $a \simeq_p b$ and $a \simeq_\gamma b$;
- a and b are incomparable w.r.t. σ^p_γ iff they are incomparable w.r.t. $\succeq_p$.

We observe that such a ranking semantics yields a total pre-order if and only if the preference relation $\succeq_p$ is a total pre-order. Otherwise, the acceptance ranking is partial. This makes sense in situations where some arguments are completely irrelevant with each other (and that explains why they are incomparable w.r.t. $\succeq_p$).

By construction, such a ranking semantics satisfies **(PP)**:

Proposition. *For every semantics γ, σ^p_γ satisfies* **(PP)**.

Roughly speaking, we first define acceptability in terms of preferences, and then we refine the sets of equally preferred arguments thanks to a “classical” ranking semantics.

To define a concrete instance of this family of semantics, we first recall the definition of h-categorizer (Besnard and Hunter 2001).

Definition 5. For $\mathcal{F} = \langle \mathcal{A}, \mathcal{R} \rangle$, we define $h : \mathcal{A} \to]0, 1]$ s.t.

$$h(a) = \frac{1}{1 + \sum_{(b,a) \in \mathcal{R}} h(b)}$$

We define the semantics based on h: $a \geq_h b$ iff $h(a) \geq h(b)$.

Now we define a preference-sensitive ranking semantics based on h-categorizer:

Definition 6. For $\mathcal{P} = \langle \mathcal{A}, \mathcal{R}, \succeq_p \rangle$, we define σ^p_h as the ranking semantics defined from $\succeq_p$ and $\geq_h$, following Definition 4.

4.2 An Application Scenario

Now we describe a scenario where our new semantics that satisfy **(PP)** can be applied. We consider the debate between candidates for an election, that is represented by arguments and attacks $\mathcal{A}$ and $\mathcal{R}$. A viewer of the debate has a preference relation $\succeq_p$ over arguments depending on the political side of the candidate who has uttered them: if two arguments have been used by candidates of the same side, they are equivalently preferred; otherwise the viewer prefers the argument that has been used by a candidate closer to his own political stand. This models the fact that the opinion of voters after viewing a debate usually fits their initial opinion (Krauss 1988; Jamieson and Birdsell 1988).

The viewer can then use one of our preference-sensitive semantics with the PAF $\mathcal{P} = \langle \mathcal{A}, \mathcal{R}, \succeq_p \rangle$. We exemplify it here with $\geq^p_h$ (Definition 6).

Suppose that the arguments exchanged by candidates A and B are these ones:

(A) a_1 = ”We should reduce the number of professors, because paying them is expensive.”

(B) a_2 = ”We cannot reduce the number of professors, actually there should be more professors since the number of students has increased recently.”

(B) a_3 = ”Moreover, on the long term, a good education system is profitable for the society and economy.”

(A) a_4 = ”This is irrelevant, there were too many professors in the past, we cannot increase their number.”

This dialog is represented by Figure 3.

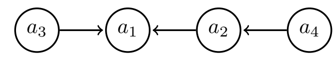

Figure 3: The Political Debate $\mathcal{P}$

We can compute the value of arguments for h-categorizer: $h(a_1) = 0.4$, $h(a_2) = 0.5$, $h(a_3) = h(a_4) = 1$.

Suppose that the debate is watched by John (who supports candidate A) and Yoko (who supports candidate B). Their preferences relations are (respectively) defined by

- $a_1 \simeq^j_p a_4 \succ^j_p a_2 \simeq^j_p a_3$;
- $a_2 \simeq^y_p a_3 \succ^y_p a_1 \simeq^y_p a_4$.

Thus, their acceptability ranking when using $\geq^p_h$ are:

- $a_4 >^{p,j}_h a_1 >^{p,j}_h a_3 >^{p,j}_h a_2$;
- $a_3 >^{p,y}_h a_2 >^{p,y}_h a_4 >^{p,y}_h a_1$.

One can observe that the acceptability of arguments reflects the viewers preferences.

5 Conclusion and Future Work

In this paper, we investigate to what extent preferences can be blended with ranking semantics for abstract argumentation frameworks. This is, to the best of our knowledge, the first study on this combination. In particular, we introduce an intuitive principle formalized by a *preference precedence* property, that describes an ideal behavior expected when rank-ordering arguments with available preferences over these arguments. We show that although existing semantics do not satisfy it, they can be used to define a new family of ranking semantics that guarantees the expected behavior. We illustrate our approach with an intuitive example which emphasizes the need for such behavior to properly model or explain human argumentation and reasoning.

This study opens some interesting perspectives and future works. While, on the first hand, it seems intuitive to apply our approach to other existing frameworks (e.g. bipolar (Cayrol and Lagasquie-Schiex 2013) or weighted argumentation (Amgoud and Doder 2018)), one can on the other hand consider other ways to combine preferences and ranking semantics. Indeed, our *preference precedence* principle aims to somehow refine provided preferences by individual acceptability associated with arguments, but a reverse approach of *preference arbitration* seems also promising.

References

Amgoud, L., and Ben-Naim, J. 2013. Ranking-based semantics for argumentation frameworks. In *Proceedings of the 7th International Conference on Scalable Uncertainty Management (SUM'13)*, 134–147.

Amgoud, L., and Cayrol, C. 2002. Inferring from inconsistency in preference-based argumentation frameworks. *J. Autom. Reasoning* 29(2):125–169.

Amgoud, L., and Doder, D. 2018. Gradual semantics for weighted graphs: An unifying approach. In *Proceedings of the Sixteenth International Conference on Principles of Knowledge Representation and Reasoning (KR'18)*, 613–614.

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

Baroni, P.; Caminada, M.; and Giacomin, M. 2018. Abstract argumentation frameworks and their semantics. In Baroni, P.; Gabbay, D.; Giacomin, M.; and van der Torre, L., eds., *Handbook of Formal Argumentation*. College Publications. 159–236.

Besnard, P., and Hunter, A. 2001. A logic-based theory of deductive arguments. *Artif. Intell.* 128(1-2):203–235.

Bonzon, E.; Delobelle, J.; Konieczny, S.; and Maudet, N. 2016. A comparative study of ranking-based semantics for abstract argumentation. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI'16)*, 914–920.

Bonzon, E.; Delobelle, J.; Konieczny, S.; and Maudet, N. 2017. A parametrized ranking-based semantics for persuasion. In *Proceedings of the 11th International Conference on Scalable Uncertainty Management (SUM'17)*, 237–251.

Cayrol, C., and Lagasquie-Schiex, M. 2005. Graduality in argumentation. *J. Artif. Intell. Res.* 23:245–297.

Cayrol, C., and Lagasquie-Schiex, M. 2013. Bipolarity in argumentation graphs: Towards a better understanding. *Int. J. Approx. Reason.* 54(7):876–899.

Cramer, M., and Guillaume, M. 2019. Empirical study on human evaluation of complex argumentation frameworks. In *Proceedings of the 16th European Conference on Logics in Artificial Intelligence (JELIA'19)*, 102–115.

Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence* 77:321–357.

Grossi, D., and Modgil, S. 2019. On the graded acceptability of arguments in abstract and instantiated argumentation. *Artif. Intell.* 275:138–173.

Jamieson, K. H., and Birdsell, D. S. 1988. *Presidential Debates: The Challenge of Creating an Informed Electorate*. Oxford University Press.

Kaci, S.; van der Torre, L. W. N.; and Villata, S. 2018. Preference in abstract argumentation. In *Proceedings of the 7th International on Computational Models of Argument (COMMA'18)*, 405–412.

Krauss, S. 1988. *Televised Presidential Debates and Public Policy*. Lawrence Erlbaum Associates.

Matt, P., and Toni, F. 2008. A game-theoretic measure of argument strength for abstract argumentation. In *Proceedings of the 11th European Conference on Logics in Artificial Intelligence (JELIA'08)*, 285–297.

Pigozzi, G.; Tsoukiàs, A.; and Viappiani, P. 2016. Preferences in artificial intelligence. *Ann. Math. Artif. Intell.* 77(3-4):361–401.

Rosenfeld, A., and Kraus, S. 2014. Argumentation theory in the field: An empirical study of fundamental notions. In *Proceedings of the Workshop on Frontiers and Connections between Argumentation Theory and Natural Language Processing, Forlì-Cesena, Italy, July 21-25, 2014*.

Rosenfeld, A., and Kraus, S. 2018. *Predicting Human Decision-Making: From Prediction to Action*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers.