

# Comparing Weak Admissibility Semantics to their Dung-style Counterparts (Extended Abstract)

Ringo Baumann<sup>1</sup>, Gerhard Brewka<sup>1</sup> and Markus Ulbricht<sup>1,2</sup>

<sup>1</sup>Department of Computer Science, Leipzig University

<sup>2</sup>Institute of Logic and Computation, TU Wien

{baumann,brewka,mulbricht}@informatik.uni-leipzig.de

## Abstract

Semantics based on weak admissibility were recently introduced to overcome a problem with self-defeating arguments that has not been solved for more than 25 years. The recursive definition of weak admissibility mainly relies on the notion of a reduct regarding a set  $E$  which only contains arguments which are neither in  $E$ , nor attacked by  $E$ . At first glance the reduct seems to be tailored for the weaker versions of Dung-style semantics only. In this paper we show that standard Dung semantics can be naturally reformulated using the reduct revealing that this concept is already implicit. We further identify a new abstract principle for semantics, so-called modularization describing how to obtain further extensions given an initial one. Its importance for the study of abstract argumentation semantics is shown by its ability to alternatively characterize classical and non-classical semantics.


## 1 Introduction

Dung's abstract argumentation frameworks (AFs) provide a formally simple basis to capture the essence of different non-monotonic formalisms [Dung, 1995]. They are focusing entirely on conflict resolution among arguments, treating the latter as abstract items without logical structure. Hence, the only information available in AFs is the so-called attack relation that determines whether an argument is in a certain conflict with another one. Coherent world views, i.e. jointly acceptable subsets of the arguments, are determined by so-called semantics.


Until recently, most of the existing argumentation semantics were either based on the concept of naivity or admissibility [van der Torre and Vesic, 2017]. The former is satisfied if the selected sets are maximal conflict-free. A set of arguments  $S$  is admissible if each attacker of an element of  $S$  is counterattacked by some element within  $S$ . This means, naive sets do not have to defend themselves against any argument whereas admissible ones have to counterattack each single attacker. In a recent paper [Baumann *et al.*, 2020b] a mediating position between these two extreme cases was introduced, so-called weak admissibility. The new concept limits the effect of self-defeating arguments, that is, arguments

which attack themselves directly or indirectly through an odd loop of arguments. Intuitively, a self-defeating argument cannot rule out an argument it attacks unless the self-defeat is eliminated by an argument breaking up the odd loop from outside. The newly introduced semantics satisfying weak admissibility are based on the notion of a reduct of an AF. Intuitively, the  $E$ -reduct of an AF is the part of the AF which is still undecided, given a conflict-free set of arguments  $E$  is accepted.

Among others, the reduct induced by a certain extension will be a central notion in the present paper. Let us consider the following example illustrating some of the core concepts we are going to investigate. Assume an agent living in Europe is planning a trip. After carefully weighing all options, the (exclusive) choice is between Paris and Las Vegas. Moreover, the agent did not yet decide whether to travel by train or airplane. Of course, Las Vegas is too far away to travel by train. The agent's knowledge base can thus be expressed by the following simple AF:



Assume the agent decides for Paris. By standard assumptions, this renders *Paris* "accepted", *Vegas* "rejected" and *train* as well as *plane* are still open. This can be formalized by the reduct of the given AF:



This reduced AF possesses both *train* and *plane* as acceptable arguments, formalizing that the agent can reach Paris both ways. With no further constraints imposed, this means both  $\{Paris, train\}$  and  $\{Paris, plane\}$  should be acceptable. The so-called modularization property will formalize this observation. If the agent decides for *Vegas*, the argument *plane* is not challenged anymore in the corresponding reduced AF, yielding  $\{Vegas, plane\}$  as unique extension.

Although these concepts appear quite natural and are indeed implicit in many AF semantics proposed in the literature, the modularization property turned out to be a surprisingly powerful tool to investigate their properties and behavior. In this extended abstract we highlight some main results around weak admissibility and modularity. In particular:

- We rephrase some of the standard notions of abstract argumentation in terms of the reduct. This sheds new light on the relationship between standard and weak admissibility semantics.
- We present new interesting properties for semantics, most notably the property of *modularization*, which go beyond the properties studied so far in abstract argumentation. These properties play a key role for the investigation of former and newly introduced semantics.

## 2 Background

Let us start by giving the necessary preliminaries.

### 2.1 Standard Concepts and Classical Semantics

We fix a non-finite background set  $\mathcal{U}$ . An argumentation framework (AF) [Dung, 1995] is a directed graph  $F = (A, R)$  where  $A \subseteq \mathcal{U}$  represents a set of arguments and  $R \subseteq A \times A$  models *attacks* between them. In this paper we consider finite AFs only (cf. [Baumann and Spanring, 2015; Baumann and Spanring, 2017] for a consideration of infinite AFs). Let  $\mathcal{F}$  denote the set of all finite AFs over  $\mathcal{U}$ . Given an AF  $F = (A, R)$  we let  $A(F) = A$  and  $R(F) = R$ . Now assume  $F = (A, R)$ . For  $U \subseteq A$  we let  $F \downarrow_U = (A \cap U, R|_{U \times U})$ . For  $a, b \in A$ , if  $(a, b) \in R$  we say that  $a$  *attacks*  $b$  as well as  $a$  *attacks* (the set)  $E$  given that  $b \in E \subseteq A$ . A set  $U \subseteq A$  is called *unattacked* if there is no  $a \in A \setminus U$  attacking  $U$ . Moreover,  $E$  is *conflict-free* in  $F$  (for short,  $E \in cf(F)$ ) iff for no  $a, b \in E$ ,  $(a, b) \in R$ . We say a set  $E$  *classically defends* (or simply,  $c$ -*defends*) an argument  $a$  if any attacker of  $a$  is attacked by some argument of  $E$ .

A *semantics*  $\sigma$  is a mapping  $\sigma : \mathcal{F} \rightarrow 2^{2^A}$  where we have  $F \mapsto \sigma(F) \subseteq 2^A$ , i.e. given an AF  $F = (A, R)$  a semantics returns a subset of  $2^A$ . In this paper we consider so-called *naive, admissible, complete, preferred, grounded* and *stable* semantics (abbr. *na, ad, co, pr, gr, stb*).

**Definition 2.1.** Let  $F = (A, R)$  be an AF and  $E \in cf(F)$ .

1.  $E \in na(F)$  iff  $E$  is  $\subseteq$ -maximal in  $cf(A)$ ,
2.  $E \in ad(F)$  iff  $E$   $c$ -defends all its elements,
3.  $E \in co(F)$  iff  $E \in ad(F)$  and for any  $x$   $c$ -defended by  $E$  we have,  $x \in E$ ,
4.  $E \in pr(F)$  iff  $E$  is  $\subseteq$ -maximal in  $co(F)$ ,
5.  $E \in gr(F)$  iff  $E$  is  $\subseteq$ -minimal in  $co(F)$ , and
6.  $E \in stb(F)$  iff  $E^+ = A \setminus E$ .

In addition, we also consider *strong admissible* sets relying on a recursive definition [Baroni and Giacomin, 2007].

**Definition 2.2.** Let  $F = (A, R)$  be an AF. A set  $E \subseteq A$  *strongly defends*  $a \in A$  if for any attacker  $b$  of  $a$ , there is some  $c \in E \setminus \{a\}$  attacking  $b$  and  $E \setminus \{a\}$  strongly defends  $c$ . A set  $E \subseteq A$  is *strongly admissible* in  $F$  ( $E \in ad^s(F)$ ) iff each  $a \in E$  is strongly defended by  $E$ .

Assume we are given an AF  $F$  and a semantics  $\sigma$ . Then we say an argument  $a \in A$  is *credulously accepted* (*skeptically accepted*) if  $a \in \bigcup \sigma(F)$  ( $a \in \bigcap \sigma(F)$ ). If  $\sigma$  is uniquely defined, i.e.  $|\sigma(F)| = 1$  for each AF  $F = (A, R)$  we may

simply speak of *accepted* arguments as both notions coincide. As usual, we slightly abuse notation and use  $\sigma \subseteq \tau$  for two semantics  $\sigma, \tau$  if  $\sigma(F) \subseteq \tau(F)$  for any AF  $F$ .


### 2.2 Reduct and Weak Admissibility

The *reduct* is a main subject of study in this paper. For a compact definition, we use  $E^+ = \{a \in A \mid E \text{ attacks } a\}$  as well as  $E^\oplus = E \cup E^+$  for a given AF  $F = (A, R)$ . The latter set is known as the *range* of  $E$ .

**Definition 2.3.** Let  $F = (A, R)$  be an AF and let  $E \subseteq A$ . The  $E$ -reduct of  $F$  is the AF  $F^E = (E^*, R \cap (E^* \times E^*))$  where  $E^* = A \setminus E^\oplus$ .

By definition,  $F^E$  is the subframework of  $F$  obtained by removing the range of  $E$  as well as corresponding attacks, i.e.  $F^E = F \downarrow_{A \setminus E^\oplus}$ . Intuitively, the  $E$ -reduct contains those arguments whose status still needs to be decided, assuming the arguments in  $E$  are accepted. Consider therefore the following illustrating example.

**Example 2.4** (Reduct and Admissibility). Let the  $F$  be the AF depicted below. In contrast to  $\{a\}$  we verify the admissibility of  $\{b\}$  in  $F$ . However, their reducts are identical and contain the self-defeating argument  $c$  only.



Observe that the reduct does not contain any attacker of the admissible set  $\{b\}$  in contrast to the non-admissible set  $\{a\}$ .

The reduct is the central notion in the definition of weak admissible semantics [Baumann *et al.*, 2020b]:

**Definition 2.5.** For an AF  $F = (A, R)$ ,  $E \subseteq A$  is called *weakly admissible* (or *w-admissible*) in  $F$  ( $E \in ad^w(F)$ ) iff

1.  $E \in cf(F)$  and
2. for any attacker  $y$  of  $E$  we have  $y \notin \bigcup ad^w(F^E)$ .

The major difference between the standard definition of admissibility and the “weak” one is that arguments do not have to  $c$ -defend themselves against *all* attackers: attackers which do not appear in any  $w$ -admissible set of the reduct can be neglected.

**Example 2.6** (Example 2.4 ctd.). In the previous example we observed  $\{a\} \notin ad(F)$ . Let us verify the weak admissibility of  $\{a\}$  in  $F$ . Obviously,  $\{a\}$  is conflict-free in  $F$  (condition 1). Moreover, since  $c$  is the only attacker of  $\{a\}$  in  $F^{\{a\}}$  we have to check  $c \notin \bigcup ad^w(F^{\{a\}})$  (condition 2). Since  $\{c\}$  violates conflict-freeness in the reduct  $F^{\{a\}} = (\{c\}, \{(c, c)\})$  we find  $\{c\} \notin ad^w(F^{\{a\}})$  yielding  $\bigcup ad^w(F^{\{a\}}) = \emptyset$ . Hence,  $c \notin \bigcup ad^w(F^{\{a\}})$  holds proving the claim.

Now *weakly preferred* semantics is defined in the natural way as  $\subseteq$ -maximal  $w$ -admissible extensions.

**Definition 2.7.** For an AF  $F = (A, R)$ ,  $E \subseteq A$  is called *weakly preferred* (or *w-preferred*) in  $F$  ( $E \in pr^w(F)$ ) iff  $E$  is  $\subseteq$ -maximal in  $ad^w(F)$ .

### 3 Semantics and their Reduct

The reduct was introduced to define weak admissibility. At first sight, it may seem that this is the only use of a somewhat ad hoc concept. However, it turns out that the notion of the reduct also helps to understand the behavior of classical AF semantics, and in particular to identify interesting connections between the classical and the new semantics. We show that classical semantics can be characterized concisely in terms of the reduct:

**Proposition 3.1.** *Let  $F = (A, R)$  be an AF and  $E \in cf(A)$ .*

1.  $E \in stb(F)$  iff  $F^E = (\emptyset, \emptyset)$ ,
2.  $E \in ad(F)$  iff no attacker of  $E$  occurs in  $F^E$ ,
3.  $E \in pr(F)$  iff  $E \in ad(F)$  and  $\bigcup ad(F^E) = \emptyset$ , and
4.  $E \in co(F)$  iff no attacker of  $E$  occurs in  $F^E$  and no argument in  $F^E$  is unattacked.

We proceed with the central modularization property. It formalizes the following intuitive idea: given a solid point of view based on an AF (an extension) and a "compatible" point of view based on the remaining AF (an extension of the reduct), these can be merged to again obtain a solid point of view (an extension of the original AF).

**Definition 3.2.** A semantics  $\sigma$  satisfies *modularization* if for any AF  $F$  we have:  $E \in \sigma(F)$  and  $E' \in \sigma(F^E)$  implies  $E \cup E' \in \sigma(F)$ .

**Proposition 3.3.** *Let  $F = (A, R)$  be an AF. Each semantics  $\sigma \in \{ad, co, pr, gr, stb\}$  satisfies modularization.*

The intuitive reason for  $\sigma = ad$  is that the reduct  $F^E$  of some  $E \in ad(F)$  i) contains only arguments that do not attack  $E$  and ii) contains all arguments that are not yet defeated by  $E$ . Hence,  $E' \in ad(F^E)$  is compatible with  $E$  and all its attackers are defeated, at least if  $E$  is present.

Thus, in some sense, one may view the modularization property as a tool to verify whether a given semantics possesses some notion of defense. It is therefore no surprise that naive semantics does not satisfy modularization. That is, a naive extension is not restrictive enough to be compatible with naive extensions of the corresponding reduct. A vanilla odd cycle suffices to illustrate this.

**Example 3.4.** Of course,  $E = \{a_1\}$  is a naive extension of  $F = (\{a_1, a_2, a_3\}, \{(a_1, a_2), (a_2, a_3), (a_3, a_1)\})$ . The corresponding reduct is  $F^E = (\{a_3\}, \emptyset)$  possessing the unique naive extension  $E' = \{a_3\}$ . Since  $E \cup E' \notin cf(F)$ , naive extensions cannot be modular.

It is easy to recognize that the previous example makes use of the fact that  $E = \{a_1\}$  does not defend itself against  $\{a_3\}$  and thus tolerates  $E' = \{a_3\}$  in the reduct  $F^E$ .

Let us return to Proposition 3.3. We claimed that preferred, grounded and stable semantics satisfy modularization. One may proof this assertion by showing that for those semantics, it must hold that  $\sigma(F^E) = \{\emptyset\}$ , whenever  $E \in \sigma(F)$ . This is stronger property than modularization and interesting on its own. We will call this property *meaningless reduct*.

**Definition 3.5.** A semantics  $\sigma$  satisfies *meaningless reduct* if for any AF  $F$  we have:  $E \in \sigma(F)$  implies  $\sigma(F^E) = \{\emptyset\}$ .

**Proposition 3.6.** *Each semantics  $\sigma \in \{pr, gr, stb\}$  satisfies meaningless reduct.*

Preferred and grounded semantics both satisfy meaningless reduct and thus also modularization. In order to distinguish them on an abstract level, we introduce a further property. We consider the following notion:

**Definition 3.7.** A semantics  $\sigma$  satisfies *unattack inclusion* if for any AF  $F$  and any unattacked argument  $a$ , there is some  $E \in \sigma(F)$  with  $a \in E$ ;  $\sigma$  satisfies *strict unattack inclusion* if for any unattacked argument  $a$ ,  $\{a\} \in \sigma(F)$  and additionally,  $\emptyset \in \sigma(F)$ .

Apart from the possibly collapsing stable semantics all classical Dung's semantics satisfy unattack inclusion. As the following Lemma formalizes, modularization even ensures that all unattacked arguments occur in the same  $\sigma$ -extension, if unattack inclusion is satisfied.

**Lemma 3.8.** *Let  $\sigma$  be any semantics satisfying modularization and unattack inclusion. If  $X$  is a set of unattacked arguments in  $F$ , then there is some  $E \in \sigma(F)$  with  $X \subseteq E$ .*

We are now in the position to characterize grounded semantics as  $\subseteq$ -least semantics regarding credulous acceptance.

**Proposition 3.9.** *For any semantics  $\sigma$  satisfying unattack inclusion and modularization we have  $\bigcup gr(F) \subseteq \bigcup \sigma(F)$  for any AF  $F$ .*

A further central result of this paper is the following: Strongly admissible semantics can be seen as the  $\subseteq$ -least semantics among all semantics satisfying strict unattack inclusion and modularization.

**Theorem 3.10.** *For any semantics  $\sigma$  satisfying strict unattack inclusion and modularization we have:  $ad^s \subseteq \sigma$ .*

### 4 Weak Admissibility Semantics


Let us now turn to the "weak" counterparts of Dung's semantics. In this section, we will discuss various properties of weak admissibility semantics. Our first observation - with a couple of interesting consequences - is that  $ad^w$  satisfies modularization as well. Since weakly admissible extensions are not admissible in general, this in particular implies that a modular semantics  $\sigma$  does not necessarily satisfy  $\sigma \subseteq ad$ .

**Theorem 4.1.** *Let  $F = (A, R)$  be an AF and  $E \in ad^w(F)$ . Suppose  $E \cap E' = \emptyset$ . Then  $E' \in ad^w(F^E)$  if and only if  $E \cup E' \in ad^w(F)$ .*


**Corollary 4.2.** *The semantics  $ad^w$  satisfies modularization.*

Let us illustrate modularization for  $ad^w$  with a slightly extended version of the AF considered in Example 2.4.

**Example 4.3.** Let  $F$  be the AF depicted below.




Since  $E = \{b\}$  is only attacked by a self-attacker in its reduct  $F^E$ ,  $E \in ad^w(F)$ . Now  $E' = \{d\} \in ad^w(F^E)$  is trivial since  $\{d\}$  is even admissible in  $F^E$ .



By modularization we obtain  $E \cup E' = \{b, d\} \in ad^w(F)$ . Indeed,  $ad^w(F) = \{\emptyset, \{b\}, \{d\}, \{b, d\}\}$ .

We want to stress that modularization also helps to restrict a given semantics  $\sigma$ . For example, the AF  $F$  cannot possess  $\{c\}$  as w-admissible extension, because modularization would enforce  $\{b, c\}$  contradicting  $ad^w(F) \subseteq cf(F)$ .



Analogously to strong admissibility (Theorem 3.10) we formulate a concise assertion regarding weak admissibility. Namely,  $ad^w$  can be seen as a  $\subseteq$ -maximal semantics among all semantics satisfying conflict-freeness and modularization. This quite surprising result can without any doubt be considered as the main theorem of this section.

**Theorem 4.4.** *For any conflict-free semantics  $\sigma$  satisfying modularization and  $ad^w \subseteq \sigma$ , we already have  $\sigma = ad^w$ .*

This means weak admissibility is among the least restrictive conflict-free semantics satisfying modularization which sheds a new light on semantics based on it. The initial motivation was to obtain a weaker version of defense, more precisely to disregard self-defeating arguments. The connection to satisfaction of the modularization property established in Theorem 4.4 is thus rather surprising: Being more liberal than  $ad^w$  already forces a semantics  $\sigma$  to either drop conflict-freeness or modularization. Moreover, it is interesting to see that strong admissible semantics is in a certain sense the most restrictive modular semantics (Theorem 3.10) while weak admissible semantics is among the most liberal ones (Theorem 4.4).

The modularization property allows us to infer that a w-preferred extension  $E$  does not tolerate existence of weakly admissible arguments in the reduct  $F^E$ . This yields a characterization of  $pr^w$  similar to classically preferred extensions, replacing preferred and admissible with w-preferred and w-admissible, respectively (see Proposition 3.1).

**Theorem 4.5.** *Let  $F = (A, R)$  be an AF. Then  $E \in pr^w(F)$  if and only if  $E$  is conflict-free such that  $\bigcup ad^w(F^E) = \emptyset$ .*

Since  $F^E$  does not possess w-admissible arguments for  $E \in pr^w(F)$ , we have  $pr^w(F^E) = \{\emptyset\}$ , implying  $pr^w$  satisfies meaningless reduct and hence also modularization.

**Corollary 4.6.** *The semantics  $pr^w$  satisfies meaningless reduct and modularization.*

Observe that we now have for any  $E \in cf(F)$

- $E \in ad^w(F)$  iff no attacker of  $E$  is in  $\bigcup ad^w(F^E)$ ,
- $E \in pr(F)$  iff  $\bigcup ad^w(F^E) = \emptyset$

which is analogous to the characterizations for the classical semantics given in Proposition 3.1, items 2 and 3.

## 5 Summary and Related Work

The investigation of argumentation semantics which rest upon weaker notions of admissibility and defense than Dung's is rather new. This is somewhat surprising as potential problems with the original versions were already pointed

out by Dung himself. In this paper we presented fundamental new results regarding weak admissibility semantics as well as classical ones. We showed that the reduct plays a key role also in the classical semantics, which sheds new light on the relationship between the new and the existing semantics. Among others, we introduced the central property of modularization playing a decisive role in finding new extensions as well as in classifying semantics. In the full version of this paper [Baumann *et al.*, 2020a] we also analyzed strong equivalence and identified the relevant kernels which allow strong equivalence to be checked by a purely syntactic transformation; an investigation of the odd cycle-free and acyclic fragments of AFs can also be found in this paper.

The handbook chapter [Baroni *et al.*, 2018] deals with modularity in AFs. It discusses and compares concepts like *directionality* and *SCC-recursiveness* [Baroni and Giacomin, 2007], *splitting* [Baumann, 2011] as well as *decomposability*, among others. The idea underlying all these concepts is the division of an AF in different parts, s.t. the semantics of the initial framework can be obtained by the semantics of the smaller parts. Such divide and conquer approaches were already successfully implemented. For instance, in [Baumann *et al.*, 2011], it was shown that splitting methods may improve the performance of algorithms computing extensions.

Weak admissibility satisfies conflict-freeness but violates classical admissibility. Conflict-tolerant semantics in contrast give up the requirement of conflict-freeness [Arieli, 2012; Grossi and Modgil, 2015]. For instance, in weighted argument systems [Dunne *et al.*, 2011] each attack is assigned a numerical weight and conflicts within extensions are allowed as long as a certain predefined inconsistency budget is not exceeded. The issue of self-defeat was already studied in [Pollock, 1987]. Pollock analyzed argument-based defeasible reasoning and proposed a semantics similar to grounded semantics. This semantics considers self-defeat as self-attack only, but not via arbitrary odd loops as we do.

This paper induces interesting future work directions. A study of the relationship between the criteria investigated in [Baroni *et al.*, 2018] and modularization would contribute to a deeper understanding of the latter; also consideration of further criteria from the literature [Amgoud and Besnard, 2013; Caminada and Amgoud, 2007]. Moreover, the capabilities of modularization when trying to characterize semantics does not appear to be exhausted at all. Finding further characterizations, maybe with the help of additional abstract criteria, is a promising future research direction. Since all semantics considered in this paper are modular in the sense of Definition 3.2, it might also be interesting to perform a more abstract and principled investigation: Why is this property implicit for so many standard AF semantics? Is modularization always connected to a certain notion of admissibility?

## Acknowledgements

This work was supported by the German Research Foundation (DFG, BA 6170/2-1), the German Federal Ministry of Education and Research (BMBF, 01/S18026A-F) by funding the competence center for Big Data and AI "ScaDS.AI" Dresden/Leipzig and by the Austrian Science Fund (FWF) through project Y698.

## References

- [Amgoud and Besnard, 2013] Leila Amgoud and Philippe Besnard. Logical limits of abstract argumentation frameworks. *Journal of Applied Non-Classical Logics*, 23(3):229–267, 2013.
- [Arieli, 2012] Ofer Arieli. Conflict-tolerant semantics for argumentation frameworks. In *Logics in Artificial Intelligence - 13th European Conference, JELIA 2012, Toulouse, France, September 26-28, 2012. Proceedings*, pages 28–40, 2012.
- [Baroni and Giacomin, 2007] Pietro Baroni and Massimiliano Giacomin. On principle-based evaluation of extension-based argumentation semantics. *Artificial Intelligence*, 171:675–700, 2007.
- [Baroni et al., 2018] Pietro Baroni, Massimiliano Giacomin, and Beishui Liao. Locality and modularity in abstract argumentation. *Handbook of Formal Argumentation*, pages 937–979, 2018.
- [Baumann and Spanring, 2015] Ringo Baumann and Christof Spanring. Infinite argumentation frameworks - on the existence and uniqueness of extensions. In *Advances in Knowledge Representation, Logic Programming, and Abstract Argumentation - Essays Dedicated to Gerhard Brewka on the Occasion of His 60th Birthday*, pages 281–295, 2015.
- [Baumann and Spanring, 2017] Ringo Baumann and Christof Spanring. A study of unrestricted abstract argumentation frameworks. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, pages 807–813, 2017.
- [Baumann et al., 2011] Ringo Baumann, Gerhard Brewka, and Renata Wong. Splitting argumentation frameworks: An empirical evaluation. In *Proceedings of the First International Workshop on the Theorie and Applications of Formal Argumentation (TAFa-11)*, pages 17–31, 2011.
- [Baumann et al., 2020a] Ringo Baumann, Gerhard Brewka, and Markus Ulbricht. Comparing weak admissibility semantics to their dung-style counterparts - reduct, modularization, and strong equivalence in abstract argumentation. In Diego Calvanese, Esra Erdem, and Michael Thielscher, editors, *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning, KR 2020, Rhodes, Greece, September 12-18, 2020*, pages 79–88, 2020.
- [Baumann et al., 2020b] Ringo Baumann, Gerhard Brewka, and Markus Ulbricht. Revisiting the foundations of abstract argumentation - semantics based on weak admissibility and weak defense. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 2742–2749, 2020.
- [Baumann, 2011] Ringo Baumann. Splitting an argumentation framework. In *Proceedings of the 11th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR-11)*, pages 40–53. Springer, 2011.
- [Caminada and Amgoud, 2007] Martin Caminada and Leila Amgoud. On the evaluation of argumentation formalisms. *Artificial Intelligence*, 171(5-6):286–310, 2007.
- [Dung, 1995] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–357, 1995.
- [Dunne et al., 2011] Paul E. Dunne, Anthony Hunter, Peter McBurney, Simon Parsons, and Michael J. Wooldridge. Weighted argument systems: Basic definitions, algorithms, and complexity results. *Artificial Intelligence*, 175(2):457–486, 2011.
- [Grossi and Modgil, 2015] Davide Grossi and Sanjay Modgil. On the graded acceptability of arguments. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 868–874, 2015.
- [Pollock, 1987] John L. Pollock. Defeasible reasoning. *Cognitive Science*, 11(4):481–518, 1987.
- [van der Torre and Vesic, 2017] Leon van der Torre and Srdjan Vesic. The principle-based approach to abstract argumentation semantics. *FLAP*, 4(8), 2017.