

On the Maximal Number of Complete Extensions in Abstract Argumentation Frameworks

Markus Ulbricht

TU Wien, Institute of Logic and Computation
Leipzig University, Department of Computer Science
mulbricht@informatik.uni-leipzig.de

Abstract

Abstract argumentation frameworks are by now a major research area in knowledge representation and reasoning. Various aspects of AFs have been extensively studied over the last 25 years. Contributing to understanding the expressive power of AFs, researchers found lower and upper bounds for the maximal number of extensions, that is, acceptable points of view, in AFs. One of the classical and most important concepts in AFs are so-called complete extensions. Surprisingly, the exact bound for the maximal number of complete extensions in an AF has not yet been formally established, although there is a reasonable conjecture tracing back at least to 2015. Recently the notion of modularization was introduced and it was shown that this concept plays a key role for the understanding of relations between semantics as well as intrinsic properties. In this paper, we will use this property to give a formal proof of the conjecture regarding complete semantics.

1 Introduction

In his seminal 1995 paper (Dung 1995), Dung initiated the investigation of abstract argumentation frameworks (AFs). Since then, various semantics have been proposed, extensions of AFs and their relationships are investigated and almost all conceivable research questions which are typical for knowledge representation formalisms are well understood. A key feature of any such formalism is its expressive power, i.e. the questions which kind of knowledge can be expressed and which not. This is not only the case for knowledge representation and reasoning, but for almost any formalism considered in theoretical computer science. Much research is driven by the expressive power of the studied framework as it hints at the need to propose extensions in order to augment its ability to model certain application scenarios. On the other hand, the search for “good” formalisms needs to manage the trade-off between expressive power and computational complexity of natural decision problems.

Both expressive power and computational complexity have been studied extensively for various extensions and semantics of AFs over the last decades. We refer the reader to (Baroni, Caminada, and Giacomin 2018) for an overview of AF semantics. The computational complexity of standard problems is well understood (Dvorák and Dunne 2018), but also more involved aspects have been studied especially in the field of dynamics like enforcing a

desired set of arguments (Wallner, Niskanen, and Järvisalo 2017; Niskanen, Wallner, and Järvisalo 2018), incorporating new beliefs (Falappa, Kern-Isberner, and Simari 2009; Haret, Wallner, and Woltran 2018) or repairing a semantical collapse (Baumann and Ulbricht 2019). More recently, the role of argumentation for explainable AI was studied (Fan and Toni 2015; Rago et al. 2020; Alfano et al. 2020; Saribatur, Wallner, and Woltran 2020; Ulbricht and Wallner 2021). Focusing on the expressive power of AF semantics, the notion of signatures was coined in (Dunne et al. 2015). In a nutshell, the signature of an AF semantics is the set of all possible sets of σ -extensions which can be associated with an AF. Even earlier, already in (Baumann and Strass 2013) a more basic question was investigated, namely how many σ -extensions can an AF possess? This question is answered for all but one classical semantics as proposed by Dung in his seminal 1995 paper. In this paper, we close this gap and confirm a conjecture mentioned in 2015 (Baumann and Strass 2015).

Interestingly, in contrast to (Baumann and Strass 2013), our results build upon recent genuine AF research (Coste-Marquis, Devred, and Marquis 2005; Baumann, Brewka, and Ulbricht 2020a) instead of graph theoretical considerations tracing back to the 1960s (Moon and Moser 1965).

2 Background

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 and we use \mathcal{F} for the set of all these graphs.

For two arguments $a, b \in A$, if $(a, b) \in R$ we say that a *attacks* b (the set E) given that $b \in E \subseteq A$. The *range* of a set E is $E^\oplus = E \cup E^+$ where $E^+ = \{a \in A \mid E \text{ attacks } a\}$. The *E -reduct* of F is the AF $F^E = (E^*, R \cap (E^* \times E^*))$ where $E^* = A \setminus E^\oplus$. We define the union $F \sqcup G$ of two AFs $F = (A, R)$ and $G = (B, S)$ as $(A \cup B, R \cup S)$. A set $E \subseteq A$ is *conflict-free* in F (for short, $E \in cf(F)$) if for no $a, b \in E$, $(a, b) \in R$. We say a set E *defends* an argument a if any attacker of a is attacked by some argument of E . A *semantics* is a function $\sigma : \mathcal{F} \rightarrow 2^{2^{\mathcal{U}}}$ with $F \mapsto \sigma(F) \subseteq 2^A$. This means, given an AF $F = (A, R)$ a semantics returns a set of subsets of A . These subsets are called *σ -extension*.

In this paper we consider so-called *admissible*, *complete*, *preferred*, and *stable* semantics (abbr. *ad*, *co*, *pr*, *stb*). They were, among others, already introduced by Dung in 1995. As usual we let $\Gamma_F(E) = \{a \in A \mid E \text{ defends } a\}$.

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

1. $E \in ad(F)$ iff $E \subseteq \Gamma_F(E)$,
2. $E \in co(F)$ iff $E = \Gamma_F(E)$,
3. $E \in pr(F)$ iff E is \subseteq -maximal in $co(F)$,
4. $E \in stb(F)$ iff $E^+ = A \setminus E$.

We recall a characterization of *co* given in (Baumann, Brewka, and Ulbricht 2020a, Proposition 3.2)

Proposition 2.2. $E \in co(F)$ iff $E \in ad(F)$ and no argument in F^E is unattacked.

3 The Conjecture


We recall the conjecture from (Baumann and Strass 2015, Conjecture 4). For a given semantics σ we let

$$\sigma_{max}(n) = \max_{F=(A,R) \in \mathcal{F}, |A| \leq n} |\sigma(F)|.$$

We have $ad_{max}(n) = 2^n$; moreover, $stb_{max}(n)$ is known (Baumann and Strass 2013) and in (Dunne et al. 2015) this result was, among others, extended to $pr_{max}(n)$. Hence, out of Dung’s classical semantics the question is answered for all semantics but *co*. The conjecture here is ($n \geq 2$)

$$co_{max}(n) = \begin{cases} 3^{n/2} & \text{if } n \text{ is even,} \\ 4 \cdot 3^{(n-3)/2} & \text{if } n \text{ is odd.} \end{cases} \quad (1)$$

This is a lower bound since the AFs indicated below (F for even, G for odd) show that these numbers can be attained.



Hence “ \geq ” in (1) was known in (Baumann and Strass 2015) already and their Conjecture 4 states that “ \leq ” is also true.

4 Preparatory Considerations

Assume we are given an AF $F = (A, R)$.

Definition 4.1. By F^{\leftrightarrow} we denote the AF F amplified by symmetric attacks, i.e. we let $F^{\leftrightarrow} = (A, R^{\leftrightarrow})$ where R^{\leftrightarrow} is given as $R^{\leftrightarrow} = \{(a, b) \mid (a, b) \in R \vee (b, a) \in R\}$

Although quite straightforward to show, the following assertion is crucial. It shows that we can restrict our investigation to symmetric AFs (Coste-Marquis, Devred, and Marquis 2005), which is a quite simple fragment.

Lemma 4.2. We have $co(F) \subseteq co(F^{\leftrightarrow})$.

Proof. Let $E \in co(F)$. We show admissibility of E in F^{\leftrightarrow} : Assume a attacks E in F^{\leftrightarrow} . By symmetry of F^{\leftrightarrow} , E counterattacks a . Now assume E defends a in F^{\leftrightarrow} . Assume E does not defend a in F . Then, there is an attacker b of a which is not attacked by E in F . Since E attacks b in F^{\leftrightarrow} , b must attack E in F since only symmetric attacks are added. This contradicts admissibility of E in F . Thus, E defends a in F , i.e. $a \in E$ and hence, $E \in co(F^{\leftrightarrow})$. \square


From now on we assume we are given a symmetric AF $F = (A, R)$. We now define the notion of a bi-cover of an AF which formalizes how to partition F into the gadget $\mathcal{G} = (\{a, b\}, \{(a, b), (b, a)\})$ “as good as possible”.¹

Definition 4.3. Let $F = (A, R)$. A *bi-cover* for F is a set $\mathcal{C} = \{C_1, \dots, C_s\}$ of subsets $C_i \subseteq A$ of A satisfying the following conditions:

- $C_i \cap C_j = \emptyset$ for $i \neq j$,
- $|C_i| = 2$ for each i ,
- $C_i = \{a_i, b_i\}$ implies $(a_i, b_i) \in R$ (and hence, due to symmetry $(b_i, a_i) \in R$ as well).

A bi-cover is *maximal* if $\bigcup C_i$ is \subseteq -maximal among all bi-covers for F . We call $\mathcal{S} = A \setminus \bigcup C_i$ the *remainder* of the cover. A cover is *exhaustive* if $\mathcal{S} = \emptyset$.

Example 4.4. Let F be the following symmetric AF which we depict as undirected graph:



A maximal bi-cover is given via $\{C_1, C_2\}$ where the sets are $C_1 = \{a, b\}$, $C_2 = \{c, d\}$. The remainder is $\{e\}$.

The remainder of a maximal bi-cover is conflict-free.

Lemma 4.5. Let \mathcal{C} be a maximal bi-cover for F and let \mathcal{S} be the remainder. Then $\mathcal{S} \in cf(F)$.

Proof. If $r_1, r_2 \in \mathcal{S}$ with $(r_1, r_2) \in R$, then $\mathcal{C} \cup \{\{r_1, r_2\}\}$ is a bi-cover with smaller remainder, contradiction. \square

In order to demonstrate the power of our bi-covers, let us mention without proof that we can already infer the conjecture for all AFs admitting an exhaustive one.

Proposition 4.6. If \mathcal{C} is an exhaustive bi-cover for F , then

$$co(F) \subseteq \{E \subseteq A \mid \forall \{a_i, b_i\} \in \mathcal{C} : a_i \notin E \vee b_i \notin E\}.$$

In particular, $|A| = n$ is even with $|co(F)| \leq 3^{n/2}$.

We require an adjustment to the modularization property (Baumann, Brewka, and Ulbricht 2020a), actually already shown in this way for weakly admissible (Baumann, Brewka, and Ulbricht 2020b) semantics.

Proposition 4.7 (Modularization Property). Assume we have $E = E' \dot{\cup} E''$ and let $E' \in ad(F)$. Then $E \in ad(F)$ if and only if $E'' \in ad(F^{E'})$.

Proof. (\Rightarrow) Clearly, $E'' \in cf(F)$. Now assume a attacks E'' in $F^{E'}$. Due to $E \in ad(F)$, some $e \in E$ counterattacks a . Since a occurs in $F^{E'}$ we infer $a \notin (E')^+$ and hence, $a \in (E'')^+$. Hence, E'' counterattacks a .

(\Leftarrow) See (Baumann, Brewka, and Ulbricht 2020a, Proposition 3.4). \square

¹This is essentially a matching in an undirected graph, but we want our terminology to focus on the structure of the arguments.

We infer the following decisive property about complete extensions:

Proposition 4.8. *Let $E \in co(F)$ with $E = E' \dot{\cup} E''$ s.t. $E' \in ad(F)$. Then $E'' \in co(F^{E'})$.*

Proof. Since all considered sets are in particular admissible, the modularization property yields $E'' \in ad(F^{E'})$. From Proposition 2.2 we know that E is complete iff $E \in ad(F)$ and F^E does not contain unattacked arguments. Since this is the case for E , $F^E = F^{E' \dot{\cup} E''} = (F^{E'})^{E''}$ does not contain unattacked arguments which means $E'' \in co(F^{E'})$. \square

Let us now return to F and a maximal bi-cover \mathcal{C} . The next step is to establish that the complete extensions of F can be computed by choosing $E' \subseteq \mathcal{S}$, considering $E'' \in co(F^{E'})$ and letting $E = E' \cup E''$.

Proposition 4.9. *If $E = E' \dot{\cup} E''$ is a complete extension for some $E' \subseteq \mathcal{S}$, then $E'' \in co(F^{E'})$. In particular, each E of this form can be constructed in such way that $E'' \subseteq \bigcup C_i$.*

Proof. As F is symmetric, $E' \subseteq \mathcal{S}$ is admissible due to being conflict-free. Hence the first claim is due to Proposition 4.8. In order to meet the condition $E'' \subseteq \bigcup C_i$, choose $E' = E \cap \mathcal{S}$. \square

5 A Combinatorial Proof

Now consider an AF F and let \mathcal{C} be a maximal bi-cover for F with remainder $\mathcal{S} = \{a_1, \dots, a_\ell\}$. For a_i , let X_i be the arguments which are linked to a_i in F .

We call each set $E' \subseteq \mathcal{S}$ a guess. By Proposition 4.9 each complete extension of F is induced by such a guess and a complete extension of the reduct $F^{E'}$. We will proceed by counting the number of guesses and the maximal number of extensions associated with it. The overall idea is that each guess excludes certain extensions due to conflicts or additional arguments being defended. Thereby, we benefit from the bi-cover as it helps us to keep track of the structure of the AF. Before proceeding to our calculations, we need to make our guesses proper: Consider such a guess $E' \subseteq \mathcal{S}$. Without loss of generality assume $E' = \{a_1, \dots, a_t\}$. If $X_i \subseteq \bigcup_{i \neq j} X_j$, then the guess E' induces the same complete extensions as $E' \setminus \{a_i\}$, because a_i is defended by the remaining arguments that have been guessed. Formally:

Definition 5.1. We call a guess $E' = \{a_1, \dots, a_t\} \subseteq \mathcal{S}$ *proper* if there is no $1 \leq i \leq t$ s.t. $X_i \subseteq \bigcup_{i \neq j} X_j$.

A simply corollary of Proposition 4.9 is thus:

Lemma 5.2. *If $E \in co(F)$, then $E = E^* \cup E''$ where $E' \subseteq \mathcal{S}$ is a proper guess, $E^* = \Gamma(E')$, and $E'' \in co(F^{E'})$.*

Proof. By Lemma 4.5, a guess which is not proper contains some argument defended by the rest of the guess. \square

We hence assume that each guess we make is proper. The following result is important for our calculations:

Lemma 5.3. *If E' is a proper guess, then the arguments in E' are in conflict with at least $|E'|$ of the sets C_i in our bi-cover.*


Proof. Let $C_i = \{c_1^i, c_2^i\}$ and assume $a \in E'$ attacks c_1^i and $b \in E'$ attacks c_2^i . Then we can move to a bi-cover with smaller remainder; contradiction. So, since our guess is proper, each argument in E' must attack some C_i which is not yet attacked by the other arguments in E' . \square

Given a proper guess $E' = \{a_1, \dots, a_t\}$, how many complete extensions of F are induced? That is, what is the size of $|\{E \in co(F) \mid E \cap \mathcal{S} = \Gamma(E')\}|$? This is a suitable starting point, because due to Lemma 5.2

$$\bigcup_{E' \text{ proper guess}} \{E \in co(F) \mid E \cap \mathcal{S} = \Gamma(E')\} = co(F) \quad (2)$$

so we may iterate over any possible proper guess. Now assume $a \in E'$, i.e. we include a into our guess. Say a attacks one (or two) arguments in $C_i = \{c_1^i, c_2^i\}$, where our bi-cover is $\{C_1, \dots, C_s\}$. If a attacks c_1^i and $a \in E'$, then our complete extensions containing E' do not contain c_1^i . Since we have no information about c_2^i , we know C_i contributes (at most) 2 choices (include c_2^i or not). As we learned from the proof of Proposition 4.6 the other sets C_j contribute at most 3 choices.

Now let $|A| = n$, $|E'| = t$ and $|\mathcal{S}| = \ell$. Our covered part contains $n - \ell$ arguments and hence the bi-cover contains $(n - \ell)/2$ sets C_i . A proper guess is in conflict with at least $|E'| = t$ sets C_i (Lemma 5.3) and is not in conflict with at most $(n - \ell)/2 - t$. Hence, the set of complete extensions E with $E \cap \mathcal{S} = \Gamma(E')$ is at most $2^t \cdot 3^{(n-\ell)/2-t}$.



Now we implement (2) by taking the sum over each possible proper guess E' , i.e. we find the upper bound

$$|co(F)| \leq \sum_{t=0}^{\ell} \binom{\ell}{t} 2^t \cdot 3^{(n-\ell)/2-t}.$$

The number n is fixed, but ℓ depends on the bi-cover. So we view this expression as a function


$$f(\ell) = \sum_{t=0}^{\ell} \binom{\ell}{t} 2^t \cdot 3^{(n-\ell)/2-t} = 3^{n/2} \cdot \left(\frac{5}{3^{3/2}}\right)^{\ell}. \quad (3)$$

This is a monotonically decreasing function in ℓ , thus attaining the maximum in $f(0) = 3^{n/2}$ which yields “ \leq ” in (1) if n is even.

For odd numbers n this is not good enough. Even if we use $\ell \geq 1$ in this case, we find a larger number than the bound $4 \cdot 3^{(n-3)/2}$ we are looking for. The reason is that our method to calculate the upper bound is too liberal as we ignore most of the structure of F .

We distinguish some cases to handle the rest of the proof.

Let us first assume that each $a \in \mathcal{S}$ possesses exactly one attack (that is one in- and one out-going attack). Suppose there are $a, a' \in \mathcal{S}$ both attacking the same set in the cover, say $C_i = \{c_1^i, c_2^i\}$. Recall the proof of Lemma 5.3: a and a' cannot attack different arguments in C_i . However, we may also w.l.o.g. assume a and a' do not attack the same argument in C_i since otherwise, $a \in E$ iff $a' \in E$ for any complete extension and we obtain the same number $|co(F)|$ as in the AF without a' . Hence, there is a bijection between \mathcal{S} and a subset of \mathcal{C} .




Now assume a attacks $c_1^i \in C_i$. If a is in our guess, then c_1^i is not, but c_2^i might or might not be (2 choices). Now comes the interesting part: If a is *not* in our guess, then c_2^i cannot be, either, because in the considered situation, c_2^i defends a . Hence we also only have 2 choices in this case. Thus, our calculation does not depend on the size of the guess E' : We have ℓ times 2 choices and exactly $(n - \ell)/2 - \ell$ sets C_i that are not attacked and thus inducing 3 choices each. We find

$$\sum_{t=0}^{\ell} \binom{\ell}{t} 2^t \cdot 3^{(n-\ell)/2-t} = 3^{3/2} \left(\frac{4}{3^{3/2}} \right)^{\ell} \cdot 3^{(n-3)/2}$$

which is again monotonically decreasing in ℓ . We assume $\ell \geq 1$ yielding $|co(F)| \leq 4 \cdot 3^{(n-3)/2}$.

Now let us consider the case where at least one a possesses two attacks. First, assume there is some C_i s.t. both arguments in C_i are attacked by $a \in \mathcal{S}$.



Recall the proof of Lemma 5.3: No argument in $\mathcal{S} \setminus \{a\}$ can attack C_i since this would induce a bi-cover with smaller remainder. Assume a is in our guess and let E'_a be the set without a . We have $(n - \ell)/2$ sets C_i , but we lose one of them due to a defeating both arguments. So if $|E'_a| = t$, we have at most $(n - \ell)/2 - t - 1$ sets C_i which are not attacked (3 choices) and t sets that still yield 2 choices. Since a is fixed, we only have $\ell - 1$ arguments to choose from, yielding at most


$$\sum_{t=0}^{\ell-1} \binom{\ell-1}{t} 2^t \cdot 3^{(n-\ell)/2-t-1} = \left(\frac{5}{3^{3/2}} \right)^{\ell-1} 3^{(n-3)/2}.$$

extensions containing a . If we do not include a , the situation is as before and we find at most additional

$$\sum_{t=0}^{\ell-1} \binom{\ell-1}{t} 2^t \cdot 3^{(n-\ell)/2-t} = 3 \cdot \left(\frac{5}{3^{3/2}} \right)^{\ell-1} 3^{(n-3)/2}.$$

extensions. Both functions are monotonically decreasing in ℓ . We add the terms and plug in $\ell = 1$ which again yields $|co(F)| \leq 4 \cdot 3^{(n-3)/2}$ as desired.

Now assume there is an argument a attacking $m \geq 2$ distinct C_i . Recall that X_i is the set of arguments in F connected to a_i . Let X be this set of arguments for a . We distinguish two cases: First assume there are $p \geq 1$ a_i s.t. $X_i \subseteq X$. This means $\ell \geq 3$ since $a \neq a_i$ and ℓ is odd.




This yields some asymmetry: A proper guess containing a does not contain these a_i , but the converse is not true. According to the m sets C_i attacked by a there are 2^m cases where a can be contained in the guess. Besides a , we have $\ell - 1 - p$ arguments to choose from and $(n - \ell)/2 - m$ sets C_i which are not yet handled. This yields at most

$$2^m \cdot \sum_{t=0}^{\ell-p-1} \binom{\ell-p-1}{t} 2^t \cdot 3^{(n-\ell)/2-t-m} \quad (4)$$

situations where a is in our guess and by standard considerations at most $\sum_{t=0}^{\ell-1} \binom{\ell-1}{t} 2^t \cdot 3^{(n-\ell)/2-t}$ where a is not. These numbers are not very convenient, but using $m \geq 2$, $p \geq 1$ and $\ell \geq 3$ yields a smaller value than $4 \cdot 3^{(n-3)/2}$.

Now assume no X_i satisfies $X_i \subseteq X$, i.e. each argument can be contained in a proper guess even if a is included.



Thus if a is included, we have $\ell - 1$ arguments to choose from. Letting $E'_a = E' \setminus \{a\}$ as before with $|E'_a| = t$, then at most $(n - \ell)/2 - t - m$ sets C_i are not attacked, yielding

$$2^m \cdot \sum_{t=0}^{\ell-1} \binom{\ell-1}{t} 2^t \cdot 3^{(n-\ell)/2-t-m}$$

proper guesses including a and by standard considerations at most $\sum_{t=0}^{\ell-1} \binom{\ell-1}{t} 2^t \cdot 3^{(n-\ell)/2-t}$ proper guesses not including a . Summing up and plugging in $\ell \geq 1$ and $m \geq 2$ yields the desired inequality $|co(F)| \leq 4 \cdot 3^{(n-3)/2}$ concluding the last case we had to consider.

6 Discussion

In this paper we verified a conjecture tracing back to 2015 by proving the precise upper bound for the maximal number of complete extensions in an AF with at most n arguments. In contrast to the technique utilized in (Baumann and Strass 2013), our proof relies on genuine recent AF research instead of finding a suitable classical graph problem and applying mathematical results. This is no surprise since defense is a concept tailored for argumentation and hence, requires techniques tailored for argumentation.

Although immediate applications of this result presumably reduce to stopping enumerating algorithms after at most $3^{n/2}$ steps, we believe this is a fundamental contribution to the foundations of AF research.

Acknowledgements

This work was supported by 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

- Alfano, G.; Calautti, M.; Greco, S.; Parisi, F.; and Trubitsyna, I. 2020. Explainable acceptance in probabilistic abstract argumentation: Complexity and approximation. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, volume 17, 33–43.
- 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.
- Baumann, R., and Strass, H. 2013. On the maximal and average numbers of stable extensions. In *International Workshop on Theorie and Applications of Formal Argumentation*, 111–126. Springer.
- Baumann, R., and Strass, H. 2015. Open problems in abstract argumentation. In *Advances in Knowledge Representation, Logic Programming, and Abstract Argumentation*. Springer. 325–339.
- Baumann, R., and Ulbricht, M. 2019. If nothing is accepted—repairing argumentation frameworks. *Journal of Artificial Intelligence Research* 66:1099–1145.
- Baumann, R.; Brewka, G.; and Ulbricht, M. 2020a. Comparing Weak Admissibility Semantics to their Dung-style Counterparts – Reduct, Modularization, and Strong Equivalence in Abstract Argumentation. In *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning*, 79–88.
- Baumann, R.; Brewka, G.; and Ulbricht, M. 2020b. Revisiting the foundations of abstract argumentation: Semantics based on weak admissibility and weak defense. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence*, 2742–2749. AAAI Press.
- Coste-Marquis, S.; Devred, C.; and Marquis, P. 2005. Symmetric argumentation frameworks. In *European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, 317–328. Springer.
- 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(2):321–357.
- Dunne, P. E.; Dvořák, W.; Linsbichler, T.; and Woltran, S. 2015. Characteristics of multiple viewpoints in abstract argumentation. *Artificial Intelligence* 228:153–178.
- Dvořák, W., and Dunne, P. E. 2018. Computational problems in formal argumentation and their complexity. In Baroni, P.; Gabbay, D.; Giacomin, M.; and van der Torre, L., eds., *Handbook of Formal Argumentation*. College Publications.
- Falappa, M. A.; Kern-Isberner, G.; and Simari, G. R. 2009. Belief revision and argumentation theory. In *Argumentation in Artificial Intelligence*. 341–360.
- Fan, X., and Toni, F. 2015. On computing explanations in argumentation. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*.
- Haret, A.; Wallner, J. P.; and Woltran, S. 2018. Two sides of the same coin: Belief revision and enforcing arguments. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, 1854–1860.
- Moon, J. W., and Moser, L. 1965. On cliques in graphs. *Israel journal of Mathematics* 3(1):23–28.
- Niskanen, A.; Wallner, J. P.; and Järvisalo, M. 2018. Extension enforcement under grounded semantics in abstract argumentation. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018, Tempe, Arizona, 30 October - 2 November 2018*, 178–183.
- Rago, A.; Cocarascu, O.; Bechlivanidis, C.; and Toni, F. 2020. Argumentation as a framework for interactive explanations for recommendations. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, volume 17, 805–815.
- Saribatur, Z. G.; Wallner, J. P.; and Woltran, S. 2020. Explaining non-acceptability in abstract argumentation. In *Proc. ECAI*, volume 325 of *Frontiers in Artificial Intelligence and Applications*, 881–888.
- Ulbricht, M., and Wallner, J. P. 2021. Strong explanations in abstract argumentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Wallner, J. P.; Niskanen, A.; and Järvisalo, M. 2017. Complexity results and algorithms for extension enforcement in abstract argumentation. *J. Artif. Intell. Res.* 60:1–40.