

Deciding Acceptance in Incomplete Argumentation Frameworks

Andreas Niskanen,¹ Daniel Neugebauer,² Matti Järvisalo,¹ Jörg Rothe²

¹ Helsinki Institute for Information Technology HIIT, Department of Computer Science, University of Helsinki, Finland

² Institut für Informatik, Heinrich-Heine-Universität Düsseldorf, Germany

Abstract

Expressing incomplete knowledge in abstract argumentation frameworks (AFs) through incomplete AFs has recently received noticeable attention. However, algorithmic aspects of deciding acceptance in incomplete AFs are still underdeveloped. We address this current shortcoming by developing algorithms for NP-hard and coNP-hard variants of acceptance problems over incomplete AFs via harnessing Boolean satisfiability (SAT) solvers. Focusing on nonempty conflict-free or admissible sets and on stable extensions, we also provide new complexity results for a refined variant of skeptical acceptance in incomplete AFs, ranging from polynomial-time computability to hardness for the second level of the polynomial hierarchy. Furthermore, central to the proposed SAT-based counterexample-guided abstraction refinement approach for the second-level problem variants, we establish conditions for redundant atomic changes to incomplete AFs from the perspective of preserving extensions. We show empirically that the resulting SAT-based approach for incomplete AFs scales at least as well as existing SAT-based approaches to deciding acceptance in AFs.

1 Introduction

Abstract argumentation frameworks (AFs), as originally proposed by Dung (1995), are today one of the most important and widely-studied formalisms in the field of argumentation in AI. Recently, a natural generalization of AFs to so-called *incomplete* argumentation frameworks (IAFs) was proposed (Coste-Marquis et al. 2007; Baumeister et al. 2018). Modeling uncertainties in terms of the existence of attacks and arguments, IAFs allow to represent structural uncertainties which may arise from dynamic changes to an AF, local views of a global AF, or uncertainties in the underlying knowledge base from which the AF is instantiated. Considering universal and existential quantification over the uncertain parts of IAFs gives rise to the respective *possible* and *necessary* generalizations of the central skeptical and credulous acceptance problems in abstract argumentation. As shown by Baumeister, Neugebauer, and Rothe (2018) via an extensive complexity analysis, these generalizations have

an impact on the computational complexity of acceptance when compared to acceptance problems in AFs.

However, in their complexity analysis, Baumeister, Neugebauer, and Rothe (2018) allowed for empty sets of acceptable arguments (extensions) or the non-existence of extensions, which may produce trivial answers to the skeptical acceptance (SA) problem for certain semantics. In particular, it would be more meaningful to restrict the semantics to nonempty extensions and to require the existence of at least one extension for a “yes” answer to SA. While SA indicates whether the target argument is among the best-accepted arguments in the AF with respect to the given semantics, the refined variant ExSA—formally defined as SA with the additional condition that an extension exists—indicates whether the target argument is actually and ultimately accepted in the IAF with respect to the semantics. To address this, we consider nonempty versions of the conflict-freeness and admissibility properties, and formally study the complexity of the variants of the more natural ExSA problem over IAFs.

Furthermore, algorithmic techniques for deciding acceptance in IAFs have not been thoroughly studied to date, and there is no clear idea of whether the recent successes in developing practical system implementations—often based on harnessing the power of Boolean satisfiability (SAT) solvers—for reasoning about acceptance in AFs can be generalized to IAFs. In this work, we address this current shortcoming by developing the first practical algorithms for deciding acceptance in IAFs.

In more detail, our contributions are the following.

- We present direct Boolean satisfiability (SAT) encodings for deciding necessary skeptical and possible credulous acceptance in IAFs for semantics under which the problems are NP- or coNP-complete (Sect. 3).
- Highlighting potential issues with allowing the non-existence of extensions in the earlier proposed acceptance problems over IAFs, we propose novel problem variants with the in-built requirement that the set of extensions must be nonempty. We provide complexity results for the new problem variants, ranging from polynomial-time to second-level completeness in the polynomial hierarchy under conflict-free, admissible, and stable semantics

(Sect. 4).

- We establish conditions for what type of atomic changes are guaranteed to be redundant from the perspective of preserving extensions of completions of IAFs (Sect. 5). While of interest on their own, this analysis proves central as a basis of SAT-based counterexample-guided abstraction refinement (CEGAR) algorithms for IAFs.
- We develop SAT-based CEGAR algorithms for the problem variants which are complete for the second-level of the polynomial hierarchy (Sect. 6).
- We empirically evaluate the scalability of the proposed SAT-based approaches to deciding acceptance in IAFs, showing that the approaches scale at least as well as state-of-the-art SAT-based approaches to deciding acceptance in (standard) AFs (Sect. 7), and that the conditions for redundant atomic changes play a central role in the scalability of the CEGAR algorithms.

2 Incomplete AFs

We formally define abstract argumentation frameworks and the evaluation semantics used in this work (an overview of other semantics is given by Baroni, Caminada, and Giacomin (2018)).

Definition 1. An *argumentation framework* (AF) consists of a (here finite) set of arguments \mathcal{A} and a binary attack relation $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$. Acceptable sets of arguments are identified via semantics, where a set $\mathcal{E} \subseteq \mathcal{A}$ is *conflict-free* (CF) if $(a, b) \notin \mathcal{R}$ for all $a, b \in \mathcal{E}$, and \mathcal{E} *defends* $a \in \mathcal{A}$ if for all $(b, a) \in \mathcal{R}$, there exists $e \in \mathcal{E}$ with $(e, b) \in \mathcal{R}$. A CF set \mathcal{E} is *admissible* (AD) if it defends all its members, and *stable* (ST) if it attacks all non-members. A set of arguments satisfying a semantics \mathbf{s} is called an \mathbf{s} extension, and we denote the set of all \mathbf{s} extensions of AF as $\mathbf{s}(AF)$.

Reasoning problems for AFs are captured by the VERIFICATION (VER) and two types of ACCEPTANCE decision problems.

Definition 2. Let $AF = \langle \mathcal{A}, \mathcal{R} \rangle$ be an AF, $\mathcal{E} \subseteq \mathcal{A}$, $a \in \mathcal{A}$, and \mathbf{s} be a semantics.

- $(AF, \mathcal{E}) \in \mathbf{s}\text{-VER}$ if and only if $\mathcal{E} \in \mathbf{s}(AF)$;
- $(AF, a) \in \mathbf{s}\text{-CREDULOUS-ACCEPTANCE}$ ($\mathbf{s}\text{-CA}$) if and only if there exists $\mathcal{E} \in \mathbf{s}(AF)$ with $a \in \mathcal{E}$; and
- $(AF, a) \in \mathbf{s}\text{-SKEPTICAL-ACCEPTANCE}$ ($\mathbf{s}\text{-SA}$) if and only if for all $\mathcal{E} \in \mathbf{s}(AF)$, it holds that $a \in \mathcal{E}$.

Since the empty set is always CF and AD, the answer to CF-SA and AD-SA is always a trivial “no.” More interesting problems are obtained when restricting both semantics to nonempty sets—we denote these restricted semantics as $\text{CF}_{\neq \emptyset}$ and $\text{AD}_{\neq \emptyset}$, respectively. Clearly, for all instances with a nonempty set of arguments, $\text{CF-VER} = \text{CF}_{\neq \emptyset}\text{-VER}$ and $\text{AD-VER} = \text{AD}_{\neq \emptyset}\text{-VER}$. Also, CA (and all its generalizations) do not distinguish between CF and $\text{CF}_{\neq \emptyset}$ or between AD and $\text{AD}_{\neq \emptyset}$.

Baumeister et al. (2018) studied incomplete AFs (IAFs) as a generalization of AFs, where the existence of distinguished arguments and attacks may be uncertain.

Definition 3. An IAF $\langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ consists of definite arguments \mathcal{A} , uncertain arguments $\mathcal{A}^?$, (conditionally) definite attacks \mathcal{R} , and uncertain attacks $\mathcal{R}^?$.

Arguments in \mathcal{A} definitely exist; arguments in $\mathcal{A}^?$ may or may not exist; attacks in \mathcal{R} exist if and only if both incident arguments exist; and attacks in $\mathcal{R}^?$ may or may not exist, but can exist only if both incident arguments exist. An IAF with $\mathcal{R}^? = \emptyset$ is called *argument-incomplete* and an IAF with $\mathcal{A}^? = \emptyset$ is called *attack-incomplete*.

Definition 4. We use $\mathcal{R}|_{\mathcal{A}'} = \mathcal{R} \cap (\mathcal{A}' \times \mathcal{A}')$ to denote the *restriction* of an attack relation to an argument set \mathcal{A}' . Any AF $\langle \mathcal{A}^*, \mathcal{R}^* \rangle$ that satisfies $\mathcal{A} \subseteq \mathcal{A}^* \subseteq (\mathcal{A} \cup \mathcal{A}^?)$ and $\mathcal{R}|_{\mathcal{A}^*} \subseteq \mathcal{R}^* \subseteq (\mathcal{R} \cup \mathcal{R}^?)|_{\mathcal{A}^*}$ is a *completion* of IAF $\langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$.

For an IAF, we say that a property defined for AFs holds *possibly* if it holds for at least one of its completions, and a property holds *necessarily* if it holds for all its completions. Accordingly, Baumeister, Neugebauer, and Rothe (2018) define possible and necessary generalizations of the CA and SA problems, which extend both problems to IAFs: $\mathbf{s}\text{-PCA}$, $\mathbf{s}\text{-PSA}$, $\mathbf{s}\text{-NCA}$, and $\mathbf{s}\text{-NSA}$.

3 SAT Encodings

We begin by providing SAT encodings for deciding acceptance in incomplete AFs for NP and coNP variants of the problem. These first encodings extend the standard encodings for argumentation semantics (Besnard and Doutre 2004) to the incomplete case by conditioning relevant parts of the formulas with the existence of an argument or an attack. These encodings also form the basis for the SAT-based counterexample-guided approach we develop in Section 6 for deciding acceptance in incomplete AFs for problems in the second level of the polynomial hierarchy.

Consider now an input IAF $IAF = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$. We use variables x_a and y_a for all $a \in \mathcal{A} \cup \mathcal{A}^?$ and $r_{a,b}$ for all $(a, b) \in \mathcal{R} \cup \mathcal{R}^?$, with the following interpretations:

- $y_a = \top$ if and only if $a \in \mathcal{A}^*$,
- $r_{a,b} = \top$ if and only if $(a, b) \in \mathcal{R}^*$, and
- $x_a = \top$ if and only if $a \in \mathcal{E} \in \mathbf{s}(AF^*)$,

where $AF^* = \langle \mathcal{A}^*, \mathcal{R}^* \rangle$ is a completion of IAF defined by the y_a and $r_{a,b}$ variables. Naturally, y_a is always true for all definite arguments, and likewise $r_{a,b}$ is true for all definite attacks. Further, if an uncertain argument is not included in the completion, it cannot be accepted, and all incident attacks are not included. This information is encoded as

$$\begin{aligned} \varphi?(IAF) = & \bigwedge_{a \in \mathcal{A}} y_a \wedge \bigwedge_{(a,b) \in \mathcal{R}} r_{a,b} \\ & \wedge \bigwedge_{a \in \mathcal{A}^?} \left(\neg y_a \rightarrow \left(\neg x_a \wedge \bigwedge_{(a,b) \in \mathcal{R}^?} \neg r_{a,b} \wedge \bigwedge_{(b,a) \in \mathcal{R}^?} \neg r_{b,a} \right) \right). \end{aligned}$$

Moving on to expressing IAF semantics, the formula

$$\varphi_{\text{CF}}(IAF) = \bigwedge_{(a,b) \in \mathcal{R} \cup \mathcal{R}^?} \left((y_a \wedge y_b \wedge r_{a,b}) \rightarrow (\neg x_a \vee \neg x_b) \right)$$

encodes conflict-free sets in a completion of IAF , expressing the fact that if two arguments and an attack between them are present in the completion, one cannot accept both of the arguments. Now, defining $z_a \leftrightarrow \bigvee_{(b,a) \in \mathcal{R} \cup \mathcal{R}^?} (x_b \wedge y_b \wedge r_{b,a})$ for each $b \in \mathcal{A} \cup \mathcal{A}^?$, the formula $\varphi_{AD}(IAF)$ defined as

$$\varphi_{CF}(IAF) \wedge \bigwedge_{a \in \mathcal{A} \cup \mathcal{A}^?} \bigwedge_{(b,a) \in \mathcal{R} \cup \mathcal{R}^?} \left((x_a \wedge y_a \wedge y_b \wedge r_{b,a}) \rightarrow z_b \right)$$

encodes admissible sets. Nonempty admissible semantics can now be encoded as $\varphi_{AD \neq \emptyset}(IAF) = \varphi_{AD}(IAF) \wedge \bigvee_{a \in \mathcal{A} \cup \mathcal{A}^?} x_a$. We express stable semantics with

$$\varphi_{ST}(IAF) = \varphi_{CF}(IAF) \wedge \bigwedge_{a \in \mathcal{A} \cup \mathcal{A}^?} \left((\neg x_a \wedge y_a) \rightarrow z_a \right),$$

setting the constraint that all arguments not included in the extension are attacked by the extension. Further, we note that complete semantics can be handled similarly via conditioning on the y_a and $r_{a,b}$ variables.

Now, for $s \in \{AD, ST\}$, the problem s -PCA is NP-complete and can be decided by the input formula

$$\varphi_{s\text{-PCA}}(IAF, a) = \varphi_s(IAF) \wedge \varphi_s(IAF) \wedge x_a$$

which is satisfiable if and only if the target argument a is possibly credulously accepted in IAF . Note that a satisfying truth assignment also yields a completion of the input IAF and an s extension of the completion where the query is included. Similarly, for $s \in \{AD \neq \emptyset, ST\}$, the problem s -NSA is coNP-complete (for $AD \neq \emptyset$ -NSA, we show this in Remark 1). The formula

$$\varphi_{s\text{-NSA}}(IAF, a) = \varphi_s(IAF) \wedge \varphi_s(IAF) \wedge \neg x_a$$

is unsatisfiable if and only if the target argument a is necessarily skeptically accepted in IAF . A satisfying truth assignment yields a counterexample completion along with an s extension not containing the query.

4 Complexity under Extension Existence

Whenever AF has no s extension, it trivially holds that $(AF, a) \in s$ -SA, even though $(AF, a) \notin s$ -CA. This behavior is exhibited by all semantics that do not guarantee the existence of an extension—in particular, the $CF \neq \emptyset$, $AD \neq \emptyset$, and ST semantics. Dunne and Wooldridge (2009) propose a refined version of SA for such semantics that, in addition, requires the existence of an extension for a “yes” answer. We call this problem EXISTENCE-AND-SKEPTICAL-ACCEPTANCE (EXSA). Formally, it is the intersection of SA and the EXISTENCE (EX) problem, or equivalently, of SA and CA. For a semantics s that guarantees $s(AF) \neq \emptyset$ for every AF , it is clear that s -EXSA = s -SA. We define generalizations of the s -EXSA problem for IAFs, namely, the s -PEXSA and s -NEXSA problems. The former is defined as follows.

s-PEXSA	
Given:	An IAF $\langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ and an argument $a \in \mathcal{A}$.
Question:	Does there exist a completion $AF^* = \langle \mathcal{A}^*, \mathcal{R}^* \rangle$ of $\langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ such that AF^* has an s extension and for each s extension \mathcal{E} of AF^* , $a \in \mathcal{E}$?

s -NEXSA is defined analogously, except that it quantifies universally (instead of existentially) over completions.

Baumeister, Neugebauer, and Rothe (2018) fully characterize the computational complexity of all variants of the CA and SA problems for IAFs. The EXSA problem and its generalizations for IAFs, as well as the restricted $CF \neq \emptyset$ and $AD \neq \emptyset$ semantics, add new open cases to this complexity landscape, which we cover in this work. We characterize the computational complexity of all variants of the skeptical acceptance problem—and in particular, of s -EXSA—for the semantics $CF \neq \emptyset$, $AD \neq \emptyset$, and ST . Our results (together with previously known results) are summarized in Table 1. We assume the reader to be familiar with the required notions from computational complexity theory, in particular, the classes of the polynomial and the Boolean hierarchy (Papadimitriou 1995; Rothe 2005), and the concepts of hardness and completeness based on polynomial-time many-one reductions. In particular, we consider deterministic and nondeterministic polynomial time, P and NP, the class DP from the second level of the Boolean hierarchy with its canonical complete problem 3-SAT-UNSAT (which is the intersection of 3-SAT and 3-UNSAT), and the classes Σ_2^P and Π_2^P from the second level of the polynomial hierarchy with their canonical complete problems Σ_2 SAT and Π_2 SAT. As noted earlier, s -EXSA differs from s -SA only for $s \in \{CF \neq \emptyset, AD \neq \emptyset, ST\}$, which is why we focus on these semantics here. Also, neither s -CA nor its generalizations distinguishes between CF and $CF \neq \emptyset$ or between AD or $AD \neq \emptyset$, so all problems based on s -CA for $s \in \{CF \neq \emptyset, AD \neq \emptyset\}$ inherit the complexity of the same problem s' -CA for the respective base semantics $s' \in \{CF, AD\}$.

Our first result is that all variants of SA for the $CF \neq \emptyset$ semantics can be efficiently decided. This is due to the fact that an argument is in some conflict-free set if and only if it has no self-attack, and it is in all conflict-free sets if and only if all other arguments have self-attacks. Uncertainty in the IAF can easily be resolved to meet these criteria.

Proposition 1. $CF \neq \emptyset$ -SA, $CF \neq \emptyset$ -EXSA, $CF \neq \emptyset$ -PSA, $CF \neq \emptyset$ -PEXSA, $CF \neq \emptyset$ -NSA, and $CF \neq \emptyset$ -NEXSA are all in P.

Next, we show DP-hardness of $AD \neq \emptyset$ -EXSA by reducing from the DP-complete problem 3-SAT-UNSAT. An illustration is given in Figure 1a. Our reduction bears some similarities to the one used by Dunne and Wooldridge (2009) in their proof that ST -EXSA is DP-hard. Our reduction is different, though, due to us having to avoid unwanted $AD \neq \emptyset$ sets that do not include the target argument.

Theorem 2. $AD \neq \emptyset$ -EXSA is DP-complete.

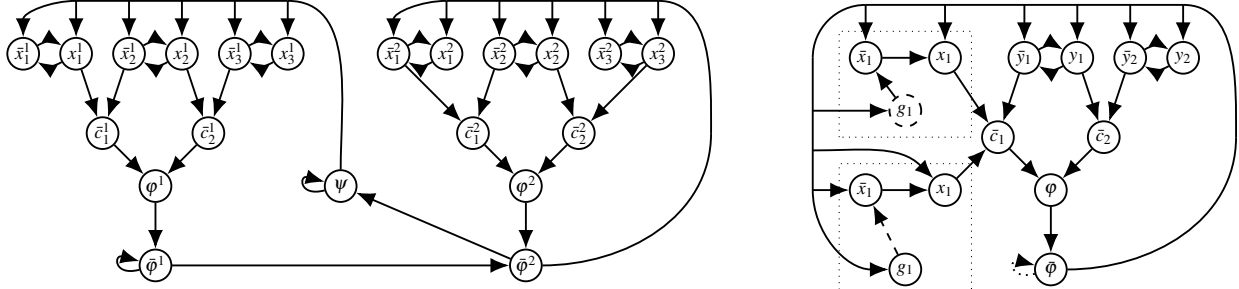
All further results are based on the standard reduction from Σ_2 SAT and Π_2 SAT problems to decision problems in IAFs, defined by Baumeister, Neugebauer, and Rothe (2018). All hardness results obtained hold even in the special cases where either $\mathcal{A}^? = \emptyset$ or $\mathcal{R}^? = \emptyset$. For ST -PEXSA, the proof of Theorem 17 by Baumeister, Neugebauer, and Rothe (2018)—which shows Σ_2^P -hardness of ST -PSA—is still valid and provides the same hardness result.

Corollary 3. ST -PEXSA is Σ_2^P -complete.

For $AD \neq \emptyset$ -PSA and $AD \neq \emptyset$ -PEXSA, we show Σ_2^P -hardness

Table 1: Overview of complexity results for variants of the standard and the refined skeptical acceptance problems, with results marked by \heartsuit due to Rey (2014); marked by \clubsuit due to Dimopoulos and Torres (1996); marked by \spadesuit due to Dunne and Wooldridge (2009); and marked by \star due to Baumeister, Neugebauer, and Rothe (2018). New results are labeled with the respective theorem or proposition number.

	s-SA		s-EXSA		s-PSA		s-PEXSA		s-NSA		s-NEXSA	
CF $\neq\emptyset$	$\in P$	(Prop. 1)	$\in P$	(Prop. 1)	$\in P$	(Prop. 1)	$\in P$	(Prop. 1)	$\in P$	(Prop. 1)	$\in P$	(Prop. 1)
AD $\neq\emptyset$	coNP-c.	\heartsuit	DP-c.	(Thm. 2)	Σ_2^p -c.	(Thm. 4)	Σ_2^p -c.	(Thm. 4)	coNP-c.	\heartsuit	Π_2^p -c.	(Thm. 5)
ST	coNP-c.	\clubsuit	DP-c.	\clubsuit	Σ_2^p -c.	\star	Σ_2^p -c.	(Cor. 3)	coNP-c.	\star	Π_2^p -c.	(Thm. 5)



(a) $(AF, \bar{\varphi}^2) \in \text{AD}_{\neq\emptyset}\text{-EXSA}$ if and only if formula φ^1 is satisfiable and formula φ^2 is unsatisfiable.

(b) $(AF, \bar{\varphi}) \in \text{AD}_{\neq\emptyset}\text{-PSA}$ if and only if there is an assignment to the X variables such that φ is unsatisfiable.

Figure 1: Illustrations of the reductions from SAT problems used in Theorem 2 (left) and Theorems 4 and 5 (right).

by extending the standard reduction with additional definite attacks from argument $\bar{\varphi}$ against all literal arguments and all grounded arguments. The reduction is illustrated in Figure 1b, where either the top framed part or the bottom framed part is included to obtain either $\mathcal{R}^? = \emptyset$ or $\mathcal{A}^? = \emptyset$. The dotted self-attack by $\bar{\varphi}$ is not included here. Our modification does not impair admissible sets that contain argument φ , but has the effect that argument $\bar{\varphi}$ is a member of all nonempty admissible sets unless formula φ is satisfiable.

Theorem 4. $\text{AD}_{\neq\emptyset}\text{-PSA}$ and $\text{AD}_{\neq\emptyset}\text{-PEXSA}$ are Σ_2^p -complete.

Remark 1. As a special case, when both $\mathcal{A}^? = \emptyset$ and $\mathcal{R}^? = \emptyset$, we obtain a reduction from 3-UNSAT to $\text{AD}_{\neq\emptyset}\text{-SA}$, which provides coNP-hardness of $\text{AD}_{\neq\emptyset}\text{-SA}$,¹ and thus also of its generalization $\text{AD}_{\neq\emptyset}\text{-NSA}$.

For $\text{AD}_{\neq\emptyset}\text{-NEXSA}$ and ST-NEXSA , we show Π_2^p -hardness by adding the definite attack $(\bar{\varphi}, \bar{\varphi})$ (dotted attack in Figure 1b). This change disables any admissible sets containing argument $\bar{\varphi}$. Here, “yes” instances of $\Pi_2^p\text{SAT}$ produce AFs that have φ in all stable sets, and “no” instances produce AFs that have completions without any nonempty admissible sets.

Theorem 5. $\text{AD}_{\neq\emptyset}\text{-NEXSA}$ and ST-NEXSA are Π_2^p -complete.

5 Preservation of Extensions

In this section, we describe atomic changes to completions (adding and removing arguments and attacks) that are guar-

anteed to be redundant from the point of view of preserving a given s extension of the completion. As we will show later on, these observations will prove to be crucial in designing efficient algorithms for the second-level problems in this work.

Let $AF^* = \langle \mathcal{A}^*, \mathcal{R}^* \rangle$ be a completion, $\mathcal{E} \in s(AF^*)$ an extension, where we denote $\text{IN}(\mathcal{E}) = \mathcal{E}$, $\text{OUT}(\mathcal{E}) = \{a \in \mathcal{A}^* \mid (b, a) \in \mathcal{R}^* \text{ with } b \in \text{IN}(\mathcal{E})\}$, and $\text{UNDEC}(\mathcal{E}) = \mathcal{A}^* \setminus (\text{IN}(\mathcal{E}) \cup \text{OUT}(\mathcal{E}))$, and let $s \in \{\text{AD}_{\neq\emptyset}, \text{ST}\}$ be a semantics.

We begin by considering adding an argument $a \in \mathcal{A}^? \setminus \mathcal{A}^*$ to the completion. If there is a definite attack $(b, a) \in \mathcal{R}$ with $b \in \text{IN}(\mathcal{E})$, then any attacks by a against arguments in \mathcal{E} would be defended by \mathcal{E} in the modified completion, which ensures that \mathcal{E} stays an extension, both under nonempty admissible and stable semantics.

Proposition 6. Let $\text{IAF} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ be an incomplete AF, $AF^* = \langle \mathcal{A}^*, \mathcal{R}^* \rangle$ a completion of IAF, $\mathcal{E} \in s(AF^*)$, and $a \in \mathcal{A}^? \setminus \mathcal{A}^*$. If there exists $(b, a) \in \mathcal{R}$ with $b \in \text{IN}(\mathcal{E})$, then $\mathcal{E} \in s(AF')$ for $AF' = \langle \mathcal{A}^* \cup \{a\}, \mathcal{R}^* \cup \mathcal{R}|_{\mathcal{A}^* \cup \{a\}} \rangle$.

We continue by considering removing arguments $a \in \mathcal{A}^? \cap \mathcal{A}^*$ from the completion. Removing arguments (along with incident attacks) that are members of $\text{OUT}(\mathcal{E})$ or of $\text{UNDEC}(\mathcal{E})$ has no effect on the current extension \mathcal{E} .

Proposition 7. Let $\text{IAF} = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ be an incomplete AF, $AF^* = \langle \mathcal{A}^*, \mathcal{R}^* \rangle$ a completion of IAF, $\mathcal{E} \in s(AF^*)$, and $a \in \mathcal{A}^? \cap \mathcal{A}^*$. If $a \notin \text{IN}(\mathcal{E})$, then $\mathcal{E} \in s(AF')$ for $AF' = \langle \mathcal{A}^* \setminus \{a\}, \mathcal{R}^*|_{\mathcal{A}^* \setminus \{a\}} \rangle$.

The preservation of a stable extension when removing and adding attacks was first studied by (Rienstra, Sakama, and van der Torre 2015); our results for stable naturally coincide, and we further extend the results to admissible sets. Remov-

¹Using different techniques, Rey (2014) showed coNP-hardness of $\text{AD}_{\neq\emptyset}\text{-SA}$.

ing an uncertain attack $(b, a) \in \mathcal{R}^? \cap \mathcal{R}^*$ with the source $b \in \text{OUT}(\mathcal{E})$ has no effect on the extension \mathcal{E} , and neither has an attack with the target $a \in \text{IN}(\mathcal{E})$. Further, the fact that \mathcal{E} is admissible is not changed by removing attacks between $\text{UNDEC}(\mathcal{E})$ arguments or from $\text{UNDEC}(\mathcal{E})$ to $\text{OUT}(\mathcal{E})$ arguments.

Proposition 8. *Let $IAF = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ be an incomplete AF, $AF^* = \langle \mathcal{A}^*, \mathcal{R}^* \rangle$ a completion of IAF , $\mathcal{E} \in \mathbf{s}(AF^*)$, and $(b, a) \in \mathcal{R}^? \cap \mathcal{R}^*$. If $b \notin \text{IN}(\mathcal{E})$ or $a \notin \text{OUT}(\mathcal{E})$, then $\mathcal{E} \in \mathbf{s}(AF')$ for $AF' = \langle \mathcal{A}^*, \mathcal{R}^* \setminus \{(b, a)\} \rangle$.*

Finally, adding an uncertain attack $(b, a) \in \mathcal{R}^? \setminus \mathcal{R}^*$ is only relevant if the source $b \notin \text{OUT}(\mathcal{E})$ and if the target $a \in \text{IN}(\mathcal{E})$, assuming $b, a \in \mathcal{A}^*$. This is due to the fact that attacks by $\text{OUT}(\mathcal{E})$ arguments are defended and thus inactivated by \mathcal{E} . Similarly, adding any attacks against $\text{OUT}(\mathcal{E})$ arguments has no effect on \mathcal{E} . Finally, the admissibility of \mathcal{E} is not changed by adding attacks between $\text{UNDEC}(\mathcal{E})$ arguments or from $\text{IN}(\mathcal{E})$ arguments to $\text{UNDEC}(\mathcal{E})$ arguments. The assumption $b, a \in \mathcal{A}^*$ suffices, since if $b \notin \mathcal{A}^*$ or $a \notin \mathcal{A}^*$, we would first need to add the argument in order to add the attack, and we only consider atomic changes.

Proposition 9. *Let $IAF = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$ be an incomplete AF, $AF^* = \langle \mathcal{A}^*, \mathcal{R}^* \rangle$ a completion of IAF , $\mathcal{E} \in \mathbf{s}(AF^*)$, and $(b, a) \in \mathcal{R}^? \setminus \mathcal{R}^*$ with $b, a \in \mathcal{A}^*$. If $b \in \text{OUT}(\mathcal{E})$ or $a \notin \text{IN}(\mathcal{E})$, then $\mathcal{E} \in \mathbf{s}(AF')$ for $AF' = \langle \mathcal{A}^*, \mathcal{R}^* \cup \{(b, a)\} \rangle$.*

We note that there is a substantial amount of research on the dynamic aspects of AFs (see (Doutre and Mailly 2018) for an overview), in particular on the impact of the addition and removal of arguments and attacks on the semantics, which we overview in Section 8.

6 Second-level Algorithms

Complementing the complexity results, we develop an iterative SAT-based approach to \mathbf{s} -PEXSA for $\mathbf{s} \in \{\text{AD}_{\neq \emptyset}, \text{ST}\}$, both of which are Σ_2^P -complete and hence beyond the reach of polynomial-size direct SAT encodings. In particular, we develop a SAT-based counterexample-guided abstraction refinement (CEGAR) approach for these problems.

In solving the problems, we are essentially looking for a witness completion in which the query argument is skeptically accepted, i.e., included in every \mathbf{s} extension. The CEGAR approach, presented in pseudocode as Algorithm 1, is based on the idea of starting with an NP abstraction of the problem at hand, overapproximating the set of witnesses for the actual problem to be solved. Concretely, we use the corresponding \mathbf{s} -PCA problem as the NP abstraction via $\text{ABSTRACTION}(IAF, a) = \varphi_{\mathbf{s}\text{-PCA}}(IAF, a)$. If $\text{ABSTRACTION}(IAF, a)$ is unsatisfiable, we can reject the query, as it is not even possibly credulously accepted. If it is satisfiable, we know that a is credulously accepted in the completion $AF^* = \text{EXTRACT}(\tau) = \langle \mathcal{A}^*, \mathcal{R}^* \rangle$ given by the satisfying truth assignment τ via

$$\begin{aligned} \mathcal{A}^* &= \{a \in \mathcal{A} \cup \mathcal{A}^? \mid \tau(y_a) = \top\}, \\ \mathcal{R}^* &= \{(a, b) \in \mathcal{R} \cup \mathcal{R}^? \mid \tau(r_{a,b}) = \top\}. \end{aligned}$$

Algorithm 1 CEGAR for possible existence and skeptical acceptance for $\mathbf{s} \in \{\text{AD}_{\neq \emptyset}, \text{ST}\}$.

Input: $IAF = \langle \mathcal{A}, \mathcal{A}^?, \mathcal{R}, \mathcal{R}^? \rangle$, $a \in \mathcal{A}$.

```

1:  $\varphi \leftarrow \text{ABSTRACTION}(IAF, a)$ 
2: while true do
3:    $(sat, \tau) \leftarrow \text{SAT}(\varphi)$ 
4:   if  $sat = \text{false}$  then return reject
5:    $AF^* \leftarrow \text{EXTRACT}(\tau)$ 
6:    $(sat, \tau) \leftarrow \text{SAT}(\text{CHECK}(IAF, AF^*, a))$ 
7:   if  $sat = \text{false}$  then return accept
8:    $\varphi \leftarrow \varphi \wedge \text{REFINE}(IAF, AF^*)$ 
9: end while

```

We continue by checking whether it is skeptically accepted, which is accomplished with the input formula

$$\begin{aligned} \text{CHECK}(IAF, AF^*, a) = \\ \varphi_{\mathbf{s}\text{-NSA}}(IAF, a) \wedge \text{COMPLETION}(IAF, AF^*), \end{aligned}$$

where $\text{COMPLETION}(IAF, AF^*)$ is defined by

$$\begin{aligned} \bigwedge_{a \in \mathcal{A}^*} y_a \wedge \bigwedge_{a \in (\mathcal{A} \cup \mathcal{A}^?) \setminus \mathcal{A}^*} \neg y_a \wedge \\ \bigwedge_{(a,b) \in \mathcal{R}^*} r_{a,b} \wedge \bigwedge_{(a,b) \in (\mathcal{R} \cup \mathcal{R}^?) \setminus \mathcal{R}^*} \neg r_{a,b} \end{aligned}$$

and encodes the current completion AF^* .

If $\text{CHECK}(IAF, AF^*, a)$ is unsatisfiable, this proves that we have successfully found a witness completion for skeptical acceptance of a , and can accept the argument. If it is satisfiable, this is not the case as we have shown that there exists an extension $\mathcal{E} \in \mathbf{s}(AF^*)$ not containing the query argument, so we refine the abstraction by adding the clause $\text{REFINE}(IAF, AF^*) = \neg \text{COMPLETION}(IAF, AF^*)$ to the SAT solver, excluding the current completion, and continue iteratively.

We would ideally like to exclude all completions which still possess the counterexample \mathcal{E} as an \mathbf{s} extension, since these would only cause more undesired iterations in the CEGAR algorithm. Due to results in Section 5, the following strong refinement is also valid, and excludes certain other completions which contain the counterexample extension:

$$\begin{aligned} \text{REFINE}(IAF, AF^*, \mathcal{E}) = \bigvee_{a \in \mathcal{A}^? \cap \text{IN}(\mathcal{E})} \neg y_a \vee \bigvee_{\substack{a \in \mathcal{A}^? \setminus \mathcal{A}^* \\ \exists (b,a) \in \mathcal{R}, b \in \text{IN}(\mathcal{E})}} y_a \\ \vee \bigvee_{(a,b) \in \mathcal{R}^? \cap \mathcal{R}^* \cap (\text{IN}(\mathcal{E}) \times \text{OUT}(\mathcal{E}))} \neg r_{a,b} \\ \vee \bigvee_{(a,b) \in (\mathcal{R}^? \setminus \mathcal{R}^*) \cap ((\text{IN}(\mathcal{E}) \cup \text{UNDEC}(\mathcal{E})) \times \text{IN}(\mathcal{E}))} r_{a,b}. \end{aligned}$$

While our focus is on the Σ_2^P -complete problems of $\text{AD}_{\neq \emptyset}$ -PEXSA and ST -PEXSA, a similar CEGAR approach can be obtained for the Π_2^P -complete problem AD -NCA by modifying Algorithm 1 as follows. We initialize the abstraction with $\varphi_{\text{AD}\text{-NSA}}(IAF, a)$, since we are now looking for

a counterexample, i.e., a completion AF^* where a is not credulously accepted. We check using $\varphi_{AD-PCA}(IAF, a) \wedge \text{COMPLETION}(IAF, AF^*)$ whether the counterexample candidate is not valid, that is, whether there is an extension containing a , implying credulous acceptance for that completion, in which case we refine the abstraction by excluding the completion. Finally, the roles of *accept* and *reject* are naturally swapped. The problem ST-NCA, which is also Π_2^P -complete, requires more modifications, since the abstraction initialized with $\varphi_{ST-NSA}(IAF, a)$ only considers those AFs which have a stable extension not containing a , that is, we would end up accepting a even though there might exist completions without a stable extension. If the abstraction becomes unsatisfiable, we iteratively call the SAT solver with input formula $\varphi_{ST-PCA}(IAF, a)$, excluding the completion using the refinement. Finally, we check using $\varphi_I(IAF) \wedge \varphi_{CF}(IAF)$ whether there still are completions; if so, we *reject* the query, since these completions have no stable extension, and otherwise we *accept* it.

7 Experiments

We continue by an overview of results from an empirical evaluation of the scalability of the approaches to acceptance problems in incomplete AFs presented in this paper. Our implementation uses Glucose 4.1 (Audemard and Simon 2018) as the underlying SAT solver, and is available in open source via <https://bitbucket.org/andreasniskanen/taeydennae>. The experiments were run on Intel Xeon E5-2680 v4 2.4-GHz, 256-GB machines with CentOS 7. We set a per-instance time limit of 900 seconds and a per-instance memory limit of 64 GB.

We generated incomplete AFs based on the ICCMA 2017 (Gaggl et al. 2016) benchmarks as follows. For each AF, we select a query argument uniformly at random from the set of arguments. Now, for each probability $p \in \{0.05, 0.1, 0.15, 0.2\}$, we generated three incomplete AFs: one where each argument (except for the query) is uncertain with probability p , one where each attack is uncertain with probability p , and one where each argument and attack is uncertain with probability p . We used the ICCMA 2017 benchmark set A for problems on the second level and the set B for problems on the first level, in-line with the complexity of the acceptance problems for which these sets were used in ICCMA 2017. This resulted in a total of 4200 IAFs for each of the two ICCMA 2017 benchmark sets.

For the NP-complete problems AD-PCA and ST-PCA, and the coNP-complete problems $AD_{\neq \emptyset}$ -NSA and ST-NSA, the mean run times (with timeouts included as the timeout limit of 900 seconds) are visualized in Figure 2, 3, and 4 (left) for different values of p and argument-incompleteness or attack-incompleteness. Interestingly, the empirical hardness of the instances does not increase as incompleteness is increased; in fact, we are able to solve the corresponding incomplete instance sets in general faster than for $p=0$, which is exactly the original ICCMA benchmark set (i.e., “normal” acceptance problems over AFs), especially when introducing attack-incompleteness. Indeed, the maximum number of timeouts (28/350) is observed on ST-NSA for $p=0$.

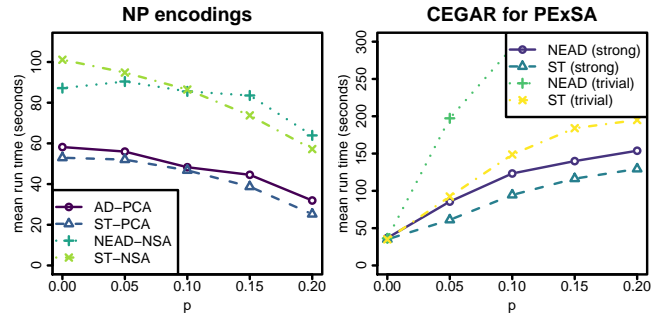


Figure 2: Mean run times (with timeouts included as 900s) for purely argument-incomplete AFs for the problems on the first (left) and the second (right) level.

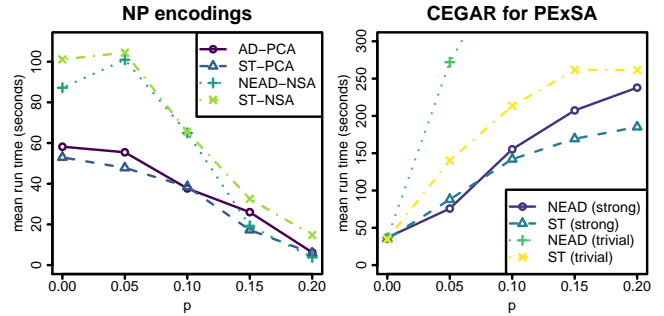


Figure 3: Mean run times (with timeouts included as 900s) for purely attack-incomplete AFs for the problems on the first (left) and the second (right) level.

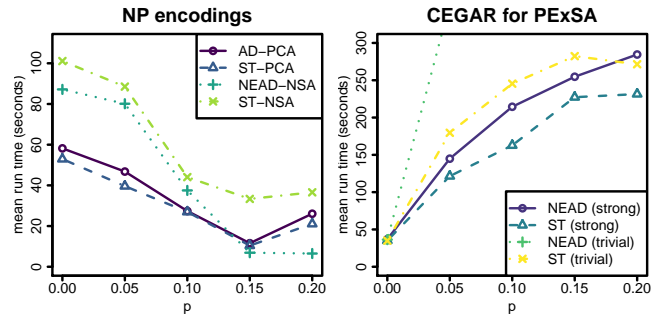


Figure 4: Mean run times (with timeouts included as 900s) for general incomplete AFs for the problems on the first (left) and the second (right) level.

The Σ_2^P -complete problems $AD_{\neq \emptyset}$ -PEXSA and ST-PEXSA are solved via the CEGAR approach. Similarly, we show the mean run times (with timeouts included) in Figure 2, 3, and 4 (right). In contrast to the NP encodings, introducing uncertainty makes the instances significantly harder to solve. We suspect this to be due to the fact that the number of potential completions to guess and check is exponential in the number of uncertain elements. However, the strong refinements, resulting from the analysis presented in Section 5, significantly speed-up the CEGAR approach; mean run times when using the trivial refinement are

considerably higher than when using the strong refinement, especially for $AD_{\neq 0}$ -PEXSA, to the extent that the strong refinements are central in making the CEGAR approach viable for deciding acceptance in incomplete AFs. As a concrete example, for $AD_{\neq 0}$ -PEXSA the average runtime over solved instances is only 24.3 seconds with $p = 0.2$ on general IAFs using the strong refinement, with 104/350 timeouts, as compared to 39.6 seconds and 225/350 timeouts for the trivial refinement, indicating that we improve especially with respect to the number of timeouts.

8 Related work

Other models that, like IAFs, aim to represent uncertainty or dynamics in abstract AFs include Partial AFs (PAFs), Control AFs (CAFs), and various types of probabilistic AFs. For PAFs (which are interchangeable with attack-incomplete AFs), Cayrol, Devred, and Lagasque-Schiex (2007) develop new, dedicated semantics, opposed to the approach of IAFs, where the semantics of an incomplete AF are derived from the standard semantics via its completions. CAFs due to Dimopoulos, Maily, and Moraitis (2018) were developed to capture strategic scenarios and, like IAFs, feature uncertain elements. Even though CAFs have a more specific role, IAF problems can be represented by corresponding CAFs problems, indicating that algorithmic and complexity results for IAFs could be reused to fuel similar results for CAFs. Finally, different types of quantitative uncertainty in AFs were proposed, where a numerical probability is associated with individual arguments or attacks (PrAFs due to Li, Oren, and Norman (2011)); with elements in an underlying knowledge base from which the AF is created (Rienstra 2012); or with entire substructures of the AF (Hunter 2014). Of these models, only PrAFs are similar to IAFs in that they allow independently uncertain arguments and attacks. However, the increased expressiveness due to the explicit probabilities—as opposed to the unquantified uncertainty in IAFs—comes at the cost of $FP^{\#P}$ -completeness for most reasoning problems in PrAFs, as shown by Fazzinga, Flesca, and Parisi (2015).

Closely related to the analysis presented in Section 5, Cayrol, Dupin de Saint-Cyr, and Lagasque-Schiex (2010) study the problem of adding an argument and incident attacks, specifically by studying necessary and sufficient conditions for satisfying different properties, also defining the atomic changes (adding and removing arguments and attacks) we study in the context of incomplete AFs. However, their focus is on grounded and preferred semantics, whereas we consider admissible and stable semantics. Semantical change when removing an argument along with incident attacks was studied by Bisquert et al. (2011) under preferred, stable, and grounded semantics, also focusing on the satisfaction of properties of extensions. The preservation of the grounded extension was studied by Boella, Kaci, and van der Torre (2009; 2010) when removing arguments and attacks, or adding attacks; again, we focus on admissible and stable extensions. Finally, the work of Rienstra, Sakama, and van der Torre (2015) focuses on the preservation of grounded, complete, preferred, stable, and semi-stable labelings under changes to the attack structure, and their results for the stable semantics coincide with ours.

The problem of adding an argument along with incident attacks is also related to expansions (Baumann and Brewka 2010; Baumann 2012b; 2012a), where sets of new arguments along with incident attacks are added to an AF. Likewise, removing arguments has been studied in the form of deletions (Baumann 2014). There may be potential for using further theoretical results of expansion, deletion, and update equivalence in order to strengthen the refinement in the CEGAR algorithm. However, equivalence is a considerably more general concept, as it concerns preserving all extensions of a given AF, whereas in our approach we are interested in preserving just the counterexample.

The SAT-based approaches developed in this work continue the successful line of work on applying SAT-based approaches to reason about argument acceptance in standard AFs (Dvořák et al. 2014; Cerutti, Giacomini, and Vallati 2019) and dynamic problems in AFs, including extension (Wallner, Niskanen, and Järvisalo 2017) and status (Niskanen, Wallner, and Järvisalo 2016) enforcement. The strong refinements we develop may also be applicable for other hard problems dealing with dynamic aspects of argumentation frameworks; so far, similar ideas have only been applied in the specific context of the NP-complete problem of extension enforcement under grounded semantics (Niskanen, Wallner, and Järvisalo 2018), in contrast to the second-level problems we develop algorithms for here.

9 Conclusion

Incomplete AFs are a natural generalization of Dung’s abstract argumentation frameworks, allowing for representing uncertainties that may arise in various settings. This work makes progress in computational aspects of incomplete AFs in terms of both theory and practice, focusing on acceptance problems. From the theoretical perspective, we proposed natural refinements of the skeptical acceptance problem and its generalizations over incomplete AFs, and provided complexity results for the new variants. From the practical perspective, we developed the first algorithmic approaches to acceptance in incomplete AF, covering both first-level (via direct SAT encodings) and second-level (via SAT-based counterexample-guided abstraction refinement) acceptance problems. Combining theory and practice, we also provided conditions for atomic changes to completions of incomplete AFs being redundant in terms of preserving an extension of a completion. This analysis is central to scaling up the SAT-based algorithms for second-level acceptance problems by allowing for significantly stronger abstraction refinement steps. The promising empirical results encourage further work on extending the SAT-based approach to acceptance in incomplete AFs to cover further semantics and problem variants; this also calls for further analysis towards strong refinements for the CEGAR approach.

Acknowledgements

This work has been financially supported in part by Academy of Finland (grants 276412, 312662, 322869), University of Helsinki Doctoral Programme in Computer Science, and by DFG grant RO 1202/14-2.

References

- Audemard, G., and Simon, L. 2018. On the Glucose SAT solver. *Int. J. Artif. Intell. Tools* 27(1):1–25.
- Baroni, P.; Caminada, M.; and Giacomin, M. 2018. Abstract argumentation frameworks and their semantics. In *Handbook of Formal Argumentation*. College Publications. 159–236.
- Baumann, R., and Brewka, G. 2010. Expanding argumentation frameworks: Enforcing and monotonicity results. In *Proc. COMMA*, volume 216 of *FAIA*, 75–86. IOS Press.
- Baumann, R. 2012a. Normal and strong expansion equivalence for argumentation frameworks. *Artif. Intell.* 193:18–44.
- Baumann, R. 2012b. What does it take to enforce an argument? Minimal change in abstract argumentation. In *Proc. ECAI*, volume 242 of *FAIA*, 127–132. IOS Press.
- Baumann, R. 2014. Context-free and context-sensitive kernels: Update and deletion equivalence in abstract argumentation. In *Proc. ECAI*, volume 263 of *FAIA*, 63–68. IOS Press.
- Baumeister, D.; Neugebauer, D.; Rothe, J.; and Schadrack, H. 2018. Verification in incomplete argumentation frameworks. *Artif. Intell.* 264:1–26.
- Baumeister, D.; Neugebauer, D.; and Rothe, J. 2018. Credulous and skeptical acceptance in incomplete argumentation frameworks. In *Proc. COMMA*, volume 305 of *FAIA*, 181–192. IOS Press.
- Besnard, P., and Doutre, S. 2004. Checking the acceptability of a set of arguments. In *Proc. NMR*, 59–64.
- Bisquert, P.; Cayrol, C.; Dupin de Saint-Cyr, F.; and Lagasquie-Schiex, M. 2011. Change in argumentation systems: Exploring the interest of removing an argument. In *Proc. SUM*, volume 6929 of *LNCS*, 275–288. Springer.
- Boella, G.; Kaci, S.; and van der Torre, L. 2009. Dynamics in argumentation with single extensions: Abstraction principles and the grounded extension. In *Proc. ECSQARU*, volume 5590 of *LNAI*, 107–118. Springer.
- Boella, G.; Kaci, S.; and van der Torre, L. 2010. Dynamics in argumentation with single extensions: Attack refinement and the grounded extension (extended version). In *Proc. ArgMAS*, volume 6057 of *LNCS*, 150–159. Springer.
- Cayrol, C.; Devred, C.; and Lagasquie-Schiex, M. 2007. Handling ignorance in argumentation: Semantics of partial argumentation frameworks. In *Proc. ECSQARU*, volume 5590 of *LNAI*, 259–270. Springer.
- Cayrol, C.; Dupin de Saint-Cyr, F.; and Lagasquie-Schiex, M. 2010. Change in abstract argumentation frameworks: Adding an argument. *J. Artif. Intell. Res.* 38:49–84.
- Cerutti, F.; Giacomin, M.; and Vallati, M. 2019. How we designed winning algorithms for abstract argumentation and which insight we attained. *Artif. Intell.* 276:1–40.
- Coste-Marquis, S.; Devred, C.; Konieczny, S.; Lagasquie-Schiex, M.; and Marquis, P. 2007. On the merging of Dung’s argumentation systems. *Artif. Intell.* 171(10):730–753.
- Dimopoulos, Y., and Torres, A. 1996. Graph theoretical structures in logic programs and default theories. *Theoret. Comput. Sci.* 170(1):209–244.
- Dimopoulos, Y.; Maily, J.; and Moraitis, P. 2018. Control argumentation frameworks. In *Proc. AAAI*, 4678–4685. AAAI Press.
- Doutre, S., and Maily, J. 2018. Constraints and changes: A survey of abstract argumentation dynamics. *Argument & Computation* 9(3):223–248.
- Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n -person games. *Artif. Intell.* 77(2):321–357.
- Dunne, P., and Wooldridge, M. 2009. Complexity of abstract argumentation. In *Argumentation in Artificial Intelligence*. Springer. 85–104.
- Dvořák, W.; Järvisalo, M.; Wallner, J. P.; and Woltran, S. 2014. Complexity-sensitive decision procedures for abstract argumentation. *Artif. Intell.* 206:53–78.
- Fazzinga, B.; Flesca, S.; and Parisi, F. 2015. On the complexity of probabilistic abstract argumentation frameworks. *ACM T. Comput. Log.* 16(3):22.
- Gaggl, S. A.; Linsbichler, T.; Maratea, M.; and Woltran, S. 2016. Introducing the second international competition on computational models of argumentation. In *Proc. SAFA*, volume 1672 of *CEUR Workshop Proceedings*, 4–9. CEUR-WS.org.
- Hunter, A. 2014. Probabilistic qualification of attack in abstract argumentation. *Int. J. Approx. Reason.* 55(2):607–638.
- Li, H.; Oren, N.; and Norman, T. 2011. Probabilistic argumentation frameworks. In *Proc. TAFE*, volume 7132 of *LNAI*, 1–16. Springer.
- Niskanen, A.; Wallner, J. P.; and Järvisalo, M. 2016. Optimal status enforcement in abstract argumentation. In *Proc. IJCAI*, 1216–1222. AAAI Press.
- Niskanen, A.; Wallner, J. P.; and Järvisalo, M. 2018. Extension enforcement under grounded semantics in abstract argumentation. In *Proc. KR*, 178–183. AAAI Press.
- Papadimitriou, C. 1995. *Computational Complexity*. Addison-Wesley, 2nd edition. Reprinted with corrections.
- Rey, L. 2014. Complexity of abstract argumentation. Master’s thesis, Heinrich-Heine-Universität Düsseldorf, Institut für Informatik, Düsseldorf, Germany.
- Rienstra, T.; Sakama, C.; and van der Torre, L. 2015. Persistence and monotony properties of argumentation semantics. In *Proc. TAFE*, volume 9524 of *LNCS*, 211–225. Springer.
- Rienstra, T. 2012. Towards a probabilistic Dung-style argumentation system. In *Proc. AT*, volume 918 of *CEUR Workshop Proceedings*, 138–152. CEUR-WS.org.
- Rothe, J. 2005. *Complexity Theory and Cryptology. An Introduction to Cryptocomplexity*. EATCS Texts in Theoretical Computer Science. Springer.
- 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.